ON SOME INFORMATION-THEORETIC ASPECTS OF NON-LINEAR STATISTICAL INVERSE PROBLEMS

RICHARD NICKL AND GABRIEL P. PATERNAIN

ABSTRACT

Results by van der Vaart (1991) from semi-parametric statistics about the existence of a non-zero Fisher information are reviewed in an infinite-dimensional non-linear Gaussian regression setting. Information-theoretically optimal inference on aspects of the unknown parameter is possible if and only if the adjoint of the linearisation of the regression map satisfies a certain range condition. It is shown that this range condition may fail in a commonly studied elliptic inverse problem with a divergence form equation ('Darcy's problem'), and that a large class of smooth linear functionals of the conductivity parameter cannot be estimated efficiently in this case. In particular, Gaussian 'Bernstein von Mises'-type approximations for Bayesian posterior distributions do not hold in this setting.

MATHEMATICS SUBJECT CLASSIFICATION 2020

Primary 62G20; Secondary 35R30

KEYWORDS

Bernstein von Mises theorems, Darcy's problem, adjoint score condition

INTERNATIONAL CONGRESS OF MATHEMATICIANS
2022 JULY 6–14

Proc. Int. Cong. Math. 2022, Vol. 7, pp. 5516–5538
DOI 10.4171/ICM2022/18

1. INTRODUCTION

The study of *inverse problems* forms an active scientific field at the interface of the physical, mathematical and statistical sciences and machine learning. A common setting is where one considers a 'forward map' $\mathscr{G}$ between two spaces of functions, and the 'inverse problem' is to recover θ from the 'data' $\mathscr{G}_\theta \equiv \mathscr{G}(\theta)$. In real-world measurement settings, data is observed *discretely*, for instance one is given point evaluations $\mathscr{G}(\theta)(X_i)$ of the function $\mathscr{G}(\theta)$ on a finite discretisation $\{X_i\}_{i=1}^N$ of the domain of $\mathscr{G}_\theta$. Each time a measurement is taken, a statistical error is incurred, and the resulting noisy data can then be described by a statistical regression model $Y_i = \mathscr{G}_\theta(X_i) + \varepsilon_i$, with regression functions $\{\mathscr{G}_\theta : \theta \in \Theta\}$ indexed by the parameter space Θ. Such models have been studied systematically at least since C. F. Gauss [9] and constitute a core part of statistical science ever since.

In a large class of important applications, the family of regression maps $\{\mathscr{G}_\theta : \theta \in \Theta\}$ arises from physical considerations and is described by a *partial differential equation* (PDE). The functional parameter θ is then naturally *infinite- (or after discretisation step, high-) dimensional*, and the map $\theta \mapsto \mathscr{G}_\theta$ is often *non-linear*, which poses challenges for statistical inference. Algorithms for such 'non-convex' problems have been proposed and developed in the last decade since influential work by A. Stuart [28], notably based on ideas from *Bayesian inference*, where the parameter θ is modelled by a Gaussian process (or related) prior Π. The inverse problem is 'solved' by approximately computing the posterior measure $\Pi(\cdot|(Y_i, X_i)_{i=1}^N)$ on Θ by an iterative (e.g. MCMC) method. While the success of this approach has become evident empirically, an objective mathematical framework that allows giving rigorous statistical and computational guarantees for such algorithms in non-linear problems has only emerged more recently. The types of results obtained so far include statistical *consistency and contraction rate* results for posterior distributions and their means, see [1,13,19] and also [14,16,21–23], as well as *computational guarantees* for MCMC based sampling schemes [3,15,25].

Perhaps the scientifically most desirable guarantees are those for 'statistical uncertainty quantification' methods based on posterior distributions, and these are notoriously difficult to obtain. Following a programme originally developed by [4–6,26] in classical 'direct' regression models, one way to address this issue is by virtue of the so-called *Bernstein–von Mises theorems* which establish asymptotically (as $N \to \infty$) exact Gaussian approximations to posterior distributions. These exploit the precise but delicate machinery from semi-parametric statistics and Le Cam theory (see [31]) and aim at showing that the actions $\langle \psi, \theta \rangle | (Y_i, X_i)_{i=1}^N$ of infinite-dimensional posterior distributions on a well-chosen set of test functions ψ converge – after rescaling by $\sqrt{N}$ (and appropriate re-centering) – to fixed normal $\mathcal{N}(0, \sigma_\theta^2(\psi))$-distributions (with high probability under the data $(Y_i, X_i)_{i=1}^N$). The limiting variance $\sigma_\theta^2(\psi)$ has an information-theoretic interpretation as the *Cramér–Rao lower bound* (inverse Fisher information) of the model (see also Section 2.4). Very few results of this type are currently available in PDE settings. Recent progress in [20] (see also related work in [12,18,21,22]) has revealed that Bernstein–von Mises theorems may hold true if the PDE underlying $\mathscr{G}_\theta$ has certain analytical properties. Specifically, one has to solve ‘information equations’ that involve the ‘information operator’ $D\mathscr{G}_\theta^* D\mathscr{G}_\theta$ generated by the linearisation $D\mathscr{G}_\theta$ of $\mathscr{G}_\theta$ (with appropriate adjoint $D\mathscr{G}_\theta^*$). The results in [20, 21] achieve this for a class of PDEs where a base differential operator (such as the Laplacian, or the geodesic vector field) is attenuated by an unknown potential θ, and where ψ can be any smooth test function.

In the present article we study a different class of elliptic PDEs commonly used to model steady state diffusion phenomena, and frequently encountered as a ‘fruitfly example’ of a non-linear inverse problem in applied mathematics (‘Darcy’s problem’; see the many references in [13, 28]). While this inverse problem can be solved in a statistically consistent way (with ‘nonparametric convergence rates’ to the ground truth, see [13, 24]), we show here that, perhaps surprisingly, semi-parametric Bernstein–von Mises phenomena for posterior distributions of a large class of linear functionals of the relevant ‘conductivity’ parameter *do in fact not hold* for this PDE, not even just locally in a ‘smooth’ neighbourhood of the standard Laplacian. See Theorems 6 and 7, which imply in particular that the inverse Fisher information $\sigma_\theta^2(\psi)$ does not exist for a large class of smooth ψ’s. The results are deduced from a theorem of van der Vaart [30] in general statistical models, combined with a thorough study of the mapping properties of $D\mathscr{G}_\theta$ and its adjoint for the PDE considered. Our negative results should help to appreciate the mathematical subtlety underpinning exact Gaussian approximations to posterior distributions in non-linear inverse problems arising with PDEs.

2. INFORMATION GEOMETRY IN NON-LINEAR REGRESSION MODELS

In this section we review some by now classical material on information-theoretical properties of infinite-dimensional regular statistical models [30, 31], and develop the details for a general vector-valued non-linear regression model relevant in inverse problems settings. Analogous results could be obtained in the idealised Gaussian white noise model (cf. Chapter 6 in [11]) sometimes considered in the inverse problems literature.

2.1. Measurement setup

Let $(\mathcal{X}, \mathcal{A}, \lambda)$ be a probability space and let V be a finite-dimensional vector space of fixed finite dimension $p_V \in \mathbb{N}$ with inner product $\langle \cdot, \cdot \rangle_V$ and norm $|\cdot|_V$. We denote by $L^\infty(\mathcal{X})$ and $L^2(\mathcal{X}) = L^2_\lambda(\mathcal{X}, V)$ the bounded measurable and λ-square integrable V-valued functions defined on $\mathcal{X}$ normed by $\|\cdot\|_\infty$ and $\|\cdot\|_{L^2_\lambda(\mathcal{X})}$, respectively. The inner product on $L^2(\mathcal{X})$ is denoted by $\langle \cdot, \cdot \rangle_{L^2(\mathcal{X})}$. We will also require Hilbert spaces $L^2(P) = L^2(V \times \mathcal{X}, P)$ of real-valued functions defined on $V \times \mathcal{X}$ that are square integrable with respect to a probability measure P on the produce space $V \times \mathcal{X}$, with inner product $\langle \cdot, \cdot \rangle_{L^2(P)}$.

We will consider a parameter space Θ that is subset of a (separable) Hilbert space $(\mathbb{H}, \langle \cdot, \cdot \rangle_{\mathbb{H}})$ on which measurable ‘forward maps’

$$\theta \mapsto \mathscr{G}(\theta) = \mathscr{G}_\theta, \quad \mathscr{G} : \Theta \to L^2_\lambda(\mathcal{X}, V), \tag{2.1}$$

are defined. Observations then arise in a general random design regression setup where one is given jointly i.i.d. random vectors $(Y_i, X_i)_{i=1}^N$ of the form

$$Y_i = \mathscr{G}_\theta(X_i) + \varepsilon_i, \quad \varepsilon_i \sim^{\text{i.i.d}} \mathcal{N}(0, I_V),\ i = 1, \dots, N, \tag{2.2}$$

where the X_i's are random i.i.d. covariates drawn from law λ on $\mathcal{X}$. We assume that the covariance I_V of each Gaussian noise vector $\varepsilon_i \in V$ is diagonal for the inner product of V. [Most of the content of this section is not specific to Gaussian errors ε_i in (2.2), cf. Example 25.28 in [31] for discussion.]

We consider a 'tangent space' H at any fixed $\theta \in \Theta$ such that H is a linear subspace of $\mathbb{H}$ and such that perturbations of θ in directions $h \in H$ satisfy $\{\theta + sh, h \in H, s \in \mathbb{R}, |s| < \epsilon\} \subset \Theta$ for some ϵ small enough. We denote by $\bar{H}$ the closure of H in $\mathbb{H}$ and will regard $\bar{H}$ itself as a Hilbert space with inner product $\langle \cdot, \cdot \rangle_{\mathbb{H}}$. We employ the following assumption in the sequel.

Condition 1. *Suppose $\mathscr{G}$ is uniformly bounded $\sup_{\theta\in\Theta} \|\mathscr{G}(\theta)\|_\infty \le U_{\mathscr{G}}$. Moreover, for fixed $\theta \in \Theta$, $x \in \mathcal{X}$, and every $h \in H$, suppose that $\mathscr{G}_\theta(x)$ is Gateaux-differentiable in direction h, that is, for all $x \in \mathcal{X}$,*

$$\left|\mathscr{G}(\theta + sh)(x) - \mathscr{G}(\theta)(x) - s\mathbb{I}_\theta[h](x)\right|_V = o(s) \quad \text{as } s \to 0, \tag{2.3}$$

for some continuous linear operator $\mathbb{I}_\theta : (H, \|\cdot\|_{\mathbb{H}}) \to L^2_\lambda(\mathcal{X}, V)$, and that for some $\epsilon > 0$ small enough and some finite constant $B = B(h, \theta)$,

$$\sup_{|s|<\epsilon} \frac{\|\mathscr{G}(\theta + sh) - \mathscr{G}(\theta)\|_\infty}{|s|} \le B. \tag{2.4}$$

2.2. The DQM property

We will now derive the semi-parametric 'score' and 'information' operators (cf. [30, 31]) in the observational model (2.2). If P_θ is the law of $(Y_1, X_1) = (\mathscr{G}(\theta)(X_1) + \varepsilon_1, X_1)$ on $V \times \mathcal{X}$ then (2.2) is an i.i.d. statistical model of product laws

$$\mathscr{P}_N = \{P_\theta^N = \otimes_{i=1}^N P_\theta : \theta \in \Theta\}, \quad N \in \mathbb{N}, \tag{2.5}$$

on $(V \times \mathcal{X})^N$, and we can identify all information-theoretic properties in terms of the model $\mathscr{P} = \mathscr{P}_1 = \{P_\theta : \theta \in \Theta\}$ for the coordinate distributions. The model $\mathscr{P}$ is *differentiable in quadratic mean (DQM)* at $\theta \in \Theta$ along the tangent space H with score operator

$$\mathbb{A}_\theta : H \to L^2(V \times \mathcal{X}, P_\theta) \tag{2.6}$$

(cf. (3.2) in [30]) if for each path $\theta_{s,h} = \theta + sh$, $h \in H$, we have as $s \to 0$,

$$\int_{V\times\mathcal{X}} \left[\frac{1}{s}\left(dP_{\theta_{s,h}}^{1/2} - dP_\theta^{1/2}\right) - \frac{1}{2}\mathbb{A}_\theta[h]\, dP_\theta^{1/2}\right]^2 \to 0 \tag{2.7}$$

where

$$dP_\theta^{1/2}(y, x) = (2\pi)^{-p_V/4} e^{-|y-\mathscr{G}(\theta)(x)|_V^2/4}\, dy\, dx, \quad y \in V, x \in \mathcal{X},$$

are the square-root probability densities of P_θ with respect to Lebesgue measure on $V \times \mathcal{X}$.

Theorem 1. *Assuming Condition 1, the model (2.5) is differentiable in quadratic mean (DQM) at $\theta \in \Theta$ along every path $(\theta + sh : |s| < \epsilon, h \in H)$ with ϵ small enough. The 'score' operator $\mathbb{A}_\theta : H \to L^2(V \times \mathcal{X}, P_\theta)$ is given by*

$$\mathbb{A}_\theta[h](y,x) = \langle y - \mathscr{G}(\theta)(x), \mathbb{I}_\theta(h)(x)\rangle_V, \quad h \in H,\ (y,x) \in V \times \mathcal{X}, \tag{2.8}$$

which extends to a continuous linear operator $\mathbb{A}_\theta : \bar{H} \to L^2(P_\theta)$.

Proof. Fix $h \in H$. Using that the densities dP_θ are strictly positive, the left-hand side in (2.7) equals

$$\int_{V\times\mathcal{X}} \left[\frac{1}{s}\left(\frac{dP_{\theta_{s,h}}^{1/2}}{dP_\theta^{1/2}} - 1\right) - \frac{1}{2}\mathbb{A}_\theta[h]\right]^2 dP_\theta$$
$$= \int_{V\times\mathcal{X}} \left[\frac{1}{s}\left[e^{\langle \frac{y}{2}, \mathscr{G}(\theta_{s,h})(x) - \mathscr{G}(\theta)(x)\rangle_V - \frac{|\mathscr{G}(\theta_{s,h})(x)|_V^2 - |\mathscr{G}(\theta)(x)|_V^2}{4}} - 1\right] - \frac{1}{2}\mathbb{A}_\theta[h]\right]^2 dP_\theta(y,x)$$
$$= \int_{V\times\mathcal{X}} \left[\frac{1}{s}\left[e^{f(s)} - 1 - \frac{s}{2}\mathbb{A}_\theta[h]\right]\right]^2 dP_\theta$$

where, for y, x fixed,

$$f(s) = \left\langle \frac{y}{2}, \mathscr{G}(\theta_{s,h})(x) - \mathscr{G}(\theta)(x)\right\rangle_V - \frac{|\mathscr{G}(\theta_{s,h})(x)|_V^2 - |\mathscr{G}(\theta)(x)|_V^2}{4}.$$

Clearly, $f(0) = 0$ and, by Condition 1 and the chain rule, we have

$$f'(0) = \left\langle \frac{y}{2}, \mathbb{I}_\theta[h](x)\right\rangle_V - \frac{\langle \mathscr{G}(\theta)(x), \mathbb{I}_\theta[h](x)\rangle_V}{2} = \frac{1}{2}\mathbb{A}_\theta[h](y,x),$$

so that the last integrand converges to zero for every $(y,x) \in V \times \mathcal{X}$, as $s \to 0$. By Condition 1 and the Cauchy–Schwarz inequality, we see that $[e^{f(s)} - 1]/s$ is bounded by a constant multiple of $e^{C|y|_V}$, $C = C(B, U_{\mathscr{G}}) < \infty$, uniformly in $|s| < \epsilon$. Furthermore, again from Condition 1,

$$\|\mathbb{A}_\theta[h]\|_{L^2(P_\theta)} \lesssim [E|Y|_V + U_{\mathscr{G}}]\|\mathbb{I}_\theta[h]\|_{L^2_\lambda} \lesssim \|h\|_{\mathbb{H}}$$

and

$$E_\theta\left[e^{C|Y|_V} + |\mathbb{A}_\theta[h](Y,X)|\right]^2 < \infty,$$

so the last limit can be P_θ-integrated by the dominated convergence theorem to give that the last displayed integral converges to zero, verifying the DQM property. The first inequality in the last display also implies that $\mathbb{A}_\theta$ extends to a continuous linear map from $\bar{H}$ to $L^2(P_\theta)$. ∎

2.3. The adjoint score and information operator

The bounded linear operator $\mathbb{A}_\theta : (\bar{H}, \langle\cdot,\cdot\rangle_{\mathbb{H}}) \to L^2(V \times \mathcal{X}, P_\theta)$ has adjoint operator

$$\mathbb{A}_\theta^* : L^2(V \times \mathcal{X}, P_\theta) \to (\bar{H}, \langle\cdot,\cdot\rangle_{\mathbb{H}})$$

which satisfies

$$\langle w, \mathbb{A}_\theta h\rangle_{L^2(P_\theta)} = \langle \mathbb{A}_\theta^* w, h\rangle_{\mathbb{H}}, \quad \text{for all } w \in L^2(V \times \mathcal{X}, P_\theta),\ h \in \bar{H}.$$

The *information operator* is then defined as

$$\mathbb{A}_\theta^* \mathbb{A}_\theta : \bar{H} \to \bar{H}. \tag{2.9}$$

Note that the 'complexity' of the statistical model enters via the choice of 'tangent space' H for which the adjoint is computed, but we suppress this in the notation.

In the present model the information operator can be entirely described in terms of the operator $\mathbb{I}_\theta : (H, \langle\cdot,\cdot\rangle_{\mathbb{H}}) \to L^2_\lambda(\mathcal{X}, V)$ from Condition 1, and its adjoint

$$\mathbb{I}_\theta^* : L^2_\lambda(\mathcal{X}, V) \to \big(\bar{H}, \langle\cdot,\cdot\rangle_{\mathbb{H}}\big).$$

Proposition 1. *Assuming Condition 1, we have*

$$\mathbb{A}_\theta^* \mathbb{A}_\theta[h] = \mathbb{I}_\theta^* \mathbb{I}_\theta[h], \quad \forall h \in H. \tag{2.10}$$

Proof. Writing ϕ for the pdf of an $\mathcal{N}(0, I_V)$ distribution, we have from Fubini's theorem, for any $w \in L^2(P_\theta)$,

$$\begin{aligned}
\langle \mathbb{A}_\theta h, w\rangle_{L^2(P_\theta)} &= \int_V \int_{\mathcal{X}} \langle y - \mathcal{G}_\theta(x), \mathbb{I}_\theta(h)(x)\rangle_V w(y,x) dP_\theta(y,x) \\
&= \int_{\mathcal{X}} \Big\langle \mathbb{I}_\theta(h)(x), \int_V \big(y - \mathcal{G}_\theta(x)\big) w(y,x) \phi\big(y - \mathcal{G}_\theta(x)\big) dy \Big\rangle_V d\lambda(x) \\
&= \big\langle \mathbb{I}_\theta(h), E_\theta\big[\big(Y - \mathcal{G}_\theta(X)\big) w(Y,X) | X = \cdot\big]\big\rangle_{L^2_\lambda} \\
&= \big\langle h, \mathbb{I}_\theta^*\big[E_\theta\big[\big(Y - \mathcal{G}_\theta(X)\big) w(Y,X) | X = \cdot\big]\big]\big\rangle_{\mathbb{H}},
\end{aligned}$$

that is, the adjoint $\mathbb{A}_\theta^* = \mathbb{I}_\theta^* \circ \mathcal{E}_\theta$ is the composition of the adjoint $\mathbb{I}_\theta^*$ of $\mathbb{I}_\theta$ with the conditional expectation (projection) operator

$$\mathcal{E}_\theta : L^2(P_\theta) \to L^2_\lambda(\mathcal{X}, V), \quad \mathcal{E}_\theta[w](x) = E_\theta\big[\big(Y - \mathcal{G}_\theta(X)\big) w(Y,X) \;\big|\; X = x\big], \quad x \in \mathcal{X}. \tag{2.11}$$

Now for $h \in H$, we see for $\varepsilon \sim \mathcal{N}(0, I_V)$ and λ-a.e. $x \in \mathcal{X}$,

$$\begin{aligned}
\mathcal{E}_\theta\big[\mathbb{A}_\theta[h]\big](x) &= E_\theta\big[\big(Y - \mathcal{G}_\theta(X)\big)\big\langle Y - \mathcal{G}_\theta(X), \mathbb{I}_\theta h(X)\big\rangle_V \;\big|\; X = x\big] \\
&= E\big[\varepsilon\langle \varepsilon, \mathbb{I}_\theta h(x)\rangle_V\big] = \mathbb{I}_\theta h(x),
\end{aligned}$$

and therefore $\mathbb{A}_\theta^* \mathbb{A}_\theta[h] = \mathbb{I}_\theta^* \mathcal{E}_\theta[\mathbb{A}_\theta[h]] = \mathbb{I}_\theta^* \mathbb{I}_\theta[h]$, completing the proof. ∎

One can think of $\mathcal{E}_\theta$ in the previous proof as a projection onto the 'space of residuals' of the regression equation (2.2), which vanishes in the representation of the information operator (2.10). In particular, the model (2.2) is LAN (locally asymptotically normal) for LAN-norm $\|\cdot\|_{\mathrm{LAN}}$ arising from LAN inner product

$$\langle h_1, h_2\rangle_{\mathrm{LAN}} := \langle \mathbb{I}_\theta h_1, \mathbb{I}_\theta h_2\rangle_{L^2_\lambda} = \langle \mathbb{A}_\theta h_1, \mathbb{A}_\theta h_2\rangle_{L^2(P_\theta)}, \quad h_1, h_2 \in \bar{H}. \tag{2.12}$$

Proposition 2. *Let $D_N \equiv (Y_i, X_i)_{i=1}^N \sim P_\theta^N$ arise from model (2.2) for some $\theta \in \Theta$ and suppose Condition 1 holds. Then the likelihood ratio process satisfies*

$$\log \frac{dP^N_{\theta + h/\sqrt{N}}}{dP_\theta}(D_N) \to^d_{N\to\infty} \mathcal{N}\Big(-\frac{1}{2}\|h\|^2_{\mathrm{LAN}}, \|h\|^2_{\mathrm{LAN}}\Big), \quad h \in H.$$

The proof follows from Theorem 1 in conjunction with Lemma 25.14 in [31] (and the central limit theorem). This, in particular, justifies the use of the terminology 'information operator' for $\mathbb{I}_\theta^*\mathbb{I}_\theta$ instead of $\mathbb{A}_\theta^*\mathbb{A}_\theta$.

In what is to follow, the range of the adjoint score operator $\mathbb{A}_\theta^*$ will play a crucial role, and we wish to record a few preparatory remarks here. By what precedes, that range equals

$$R(\mathbb{A}_\theta^*) = \{\psi = \mathbb{I}_\theta^* \mathscr{E}_\theta w, \text{ for some } w \in L^2(P_\theta)\}, \tag{2.13}$$

where $\mathscr{E}_\theta$ is from (2.11). Since $\mathscr{E}_\theta$ maps $L^2(P_\theta)$ into L^2_λ, a fortiori any $\psi \in R(\mathbb{A}_\theta^*)$ has to satisfy

$$\psi \in R(\mathbb{I}_\theta^*) = \{\psi = \mathbb{I}_\theta^* h, \text{ for some } h \in L^2_\lambda(\mathcal{X}, V)\}, \tag{2.14}$$

so $R(\mathbb{A}_\theta^*) \subset R(\mathbb{I}_\theta^*)$. Likewise, taking $w(y,x) = \langle y - \mathscr{G}(\theta)(x), h(x)\rangle_V \in L^2(P_\theta)$, we can realise (arguing as in the proof of the last proposition) any $h \in L^2_\lambda(\mathcal{X})$ as $\mathscr{E}_\theta w = h$ and so if $\psi \in R(\mathbb{I}_\theta^*)$ then $\psi \in R(\mathbb{A}_\theta^*)$, too. We conclude that

$$R(\mathbb{I}_\theta^*) = R(\mathbb{A}_\theta^*). \tag{2.15}$$

2.4. Lower bounds for estimation of functionals

Suppose the problem is to estimate a *linear* functional $\Psi : \Theta \to \mathbb{R}$ of the unknown parameter θ. Let

$$\mathscr{P}_H := \{w = \mathbb{A}_\theta(h) : h \in H\} \subset L^2(V \times \mathcal{X}, P_\theta)$$

denote the tangent space of the model $\mathscr{P}$ induced by H. Suppose further we can find $\tilde{\psi}_\theta \in L^2(P_\theta)$ (the 'efficient influence function') such that

$$\Psi(h) = \langle \tilde{\psi}_\theta, \mathbb{A}_\theta h\rangle_{L^2(P_\theta)}, \quad h \in H. \tag{2.16}$$

If such $\tilde{\psi}_\theta$ exists, we can always take it to belong to the closure $\overline{\mathscr{P}_H}$ of $\mathscr{P}_H$ in $L^2(P_\theta)$ (simply by $L^2(P_\theta)$-projection onto $\overline{\mathscr{P}_H}$, if necessary). A lower bound for the optimal efficient asymptotic variance for $\sqrt{N}$-consistent estimators of $\Psi(\theta)$ over the model $\{\theta + h/\sqrt{N}, h \in H\}$ is then given by

$$\sup_{0 \neq w \in \mathscr{P}_H} \frac{\langle \tilde{\psi}_\theta, w\rangle^2_{L^2(P_\theta)}}{\langle w, w\rangle_{L^2(P_\theta)}} = \|\tilde{\psi}_\theta\|^2_{L^2(P_\theta)}, \tag{2.17}$$

with equality holding in view of $\tilde{\psi}_\theta \in \overline{\mathscr{P}_H}$ and the Cauchy–Schwarz inequality. Specifically, by Theorem 25.21 in [31], one has

$$\liminf_{N\to\infty} \inf_{\tilde{\psi}_N : (V\times\mathcal{X})^N \to \mathbb{R}} \sup_{h \in H, \|h\|_{\mathbb{H}} \le 1/\sqrt{N}} N E^N_{\theta+h} \big(\tilde{\psi}_N - \Psi(\theta + h)\big)^2 \ge \|\tilde{\psi}_\theta\|^2_{L^2(P_\theta)}. \tag{2.18}$$

If the functional is of the form $\Psi(h) = \langle \psi, h\rangle_{\mathbb{H}}$ for some fixed test function ψ, and if $\mathbb{A}_\theta^*$ is the adjoint of $\mathbb{A}_\theta$ from the previous subsection, the requirement (2.16) can be written as

$$\langle \psi, h\rangle_{\mathbb{H}} = \langle \tilde{\psi}_\theta, \mathbb{A}_\theta h\rangle_{L^2(P_\theta)} = \langle \mathbb{A}_\theta^* \tilde{\psi}_\theta, h\rangle_{\mathbb{H}}, \quad h \in H, \tag{2.19}$$

and hence reduces to $\psi = \mathbb{A}_\theta^* \tilde{\psi}_\theta$ for some $\tilde{\psi}_\theta \in L^2(P_\theta)$, that is, $\psi \in R(\mathbb{A}_\theta^*)$ from (2.13).

2.5. Non-existence of $\sqrt{N}$-consistent estimators of linear functionals

Arguing along the traditional lines of the proof of the Cramer–Rao inequality, the inverse of

$$i_{\theta,h,\psi} := \frac{\|\mathbb{A}_\theta h\|^2_{L^2(P_\theta)}}{\langle \psi, h\rangle^2_{\mathbb{H}}} \tag{2.20}$$

provides an a priori lower bound for the variance of any estimator $\widehat{\Psi}$ of $\Psi(\theta) = \langle \psi, \theta\rangle_{\mathbb{H}}$ that is unbiased (i.e. satisfies $E_\theta \widehat{\Psi} = \Psi(\theta)$) for all θ in the one-dimensional model $\{\theta + sh : |s| < \epsilon\}$. The *efficient* Fisher information for estimating Ψ optimally for all elements $h \in H$ of the tangent space is then given by

$$i_{\theta,H,\psi} := \inf_{h\in H, \langle\psi,h\rangle_{\mathbb{H}}\neq 0} \frac{\|\mathbb{A}_\theta h\|^2_{L^2(P_\theta)}}{\langle \psi, h\rangle^2_{\mathbb{H}}}. \tag{2.21}$$

Note that when $\psi = \mathbb{A}^*_\theta \tilde{\psi}_\theta$ is in the range of $\mathbb{A}^*_\theta$ then we can rewrite the last number as

$$\inf_{h\in H, \langle\psi,h\rangle_{\mathbb{H}}\neq 0} \frac{\|\mathbb{A}_\theta h\|^2_{L^2(P_\theta)}}{\langle \mathbb{A}^*_\theta \tilde{\psi}_\theta, h\rangle^2_{\mathbb{H}}} = \inf_{h\in H, \langle\psi,h\rangle_{\mathbb{H}}\neq 0} \frac{\|\mathbb{A}_\theta h\|^2_{L^2(P_\theta)}}{\langle \tilde{\psi}_\theta, \mathbb{A}_\theta h\rangle^2_{L^2(P_\theta)}}. \tag{2.22}$$

Since $\psi \in R(\mathbb{A}^*_\theta)$ is orthogonal on $\ker(\mathbb{A}_\theta)$, using also (2.17), we thus arrive at

$$\|\tilde{\psi}_\theta\|^2_{L^2(P_\theta)} = \sup_{h\in H, \mathbb{A}_\theta h\neq 0} \frac{\langle \tilde{\psi}_\theta, \mathbb{A}_\theta h\rangle^2_{L^2(P_\theta)}}{\langle \mathbb{A}_\theta h, \mathbb{A}_\theta h\rangle_{L^2(P_\theta)}} = i^{-1}_{\theta,H,\psi}, \tag{2.23}$$

explaining the relationship to the best asymptotic variance in (2.18).

An important observation of van der Vaart (Theorem 4.1 in [30]) is that a necessary and sufficient condition for the Fisher information for estimating $\Psi(\theta) = \langle \theta, \psi\rangle_{\mathbb{H}}$ to be non-zero is that ψ indeed lies in the range of $\mathbb{A}^*_\theta$.

Theorem 2. *For $\theta \in \Theta$ and tangent space H, let $i_{\theta,H,\psi}$ be the efficient Fisher information (2.21) for estimating the functional $\Psi(\theta) = \langle \theta, \psi\rangle_{\mathbb{H}}$, $\psi \in \bar{H}$. Then $i_{\theta,H,\psi} > 0$ if and only if $\psi \in R(\mathbb{I}^*_\theta)$.*

If $\psi \in R(\mathbb{I}^*_\theta)$ then positivity $i_{\theta,H,\psi} > 0$ follows directly from (2.15), (2.22) and the Cauchy–Schwarz inequality. The converse is slightly more involved – we include a proof in Section 4.2 below for the case most relevant in inverse problems when the information operator $\mathbb{I}^*_\theta \mathbb{I}_\theta$ from (2.10) is *compact* on $\bar{H}$ (see after Proposition 4 below for the example relevant here).

It follows that if $\psi \notin R(\mathbb{I}^*_\theta)$ then $\Psi(\theta)$ cannot be estimated at $\sqrt{N}$-rate.

Theorem 3. *Consider estimating a functional $\Psi(\theta) = \langle \psi, \theta\rangle_{\mathbb{H}}$, $\psi \in \bar{H}$, based on i.i.d. data $(Y_i, X_i)_{i=1}^N$ in the model (2.2) satisfying Condition 1 for some $\theta \in \Theta$ and tangent space H. Suppose $i_{\theta,H,\psi} = 0$. Then*

$$\liminf_{N\to\infty} \inf_{\tilde{\psi}_N:(V\times\mathcal{X})^N\to\mathbb{R}} \sup_{h\in H, \|h\|_{\mathbb{H}}\leq 1/\sqrt{N}} N E^N_{\theta+h}\big(\tilde{\psi}_N - \Psi(\theta+h)\big)^2 = \infty. \tag{2.24}$$

The last theorem can be proved following the asymptotic arguments leading to the proof of (2.18) in Theorem 25.21 in [31]. A proof that follows more directly from the preceding developments is as follows: Augment the observation space to include measurements

$(Z_i, Y_i, X_i)_{i=1}^N \sim \bar{P}_\theta^N$ where the $Z_i \sim^{iid} \mathcal{N}(\langle \theta, \psi \rangle_{\mathbb{H}}, \sigma^2)$ are independent of the (Y_i, X_i)'s, and where σ^2 is known but arbitrary. The new model $\bar{\mathscr{P}}_N = \{\bar{P}_\theta^N : \theta \in \Theta\}$ has 'augmented' LAN norm from (2.12) given by

$$\|\bar{\mathbb{A}}_\theta h\|^2_{L^2(\bar{P}_\theta)} = \|\mathbb{A}_\theta h\|^2_{L^2(P_\theta)} + \sigma^{-2} \langle \psi, h \rangle^2_{\mathbb{H}}, \quad h \in \bar{H},$$

as can be seen from a standard tensorisation argument for independent sample spaces and the fact that a $\mathcal{N}(\langle \theta, \psi \rangle_{\mathbb{H}}, \sigma^2)$ model has LAN 'norm' $\sigma^{-2} \langle \psi, h \rangle^2_{\mathbb{H}}$, by a direct calculation with Gaussian densities. In particular, the efficient Fisher information from (2.21) for estimating $\langle \psi, \theta \rangle_{\mathbb{H}}$ from the augmented data is now of the form

$$\bar{i}_{\theta,H,\psi} = \inf_h \frac{\|\mathbb{A}_\theta h\|^2_{L^2(P_\theta)} + \sigma^{-2} \langle \psi, h \rangle^2_{\mathbb{H}}}{\langle \psi, h \rangle^2_{\mathbb{H}}} = i_{\theta,H,\psi} + \sigma^{-2} = \sigma^{-2} > 0.$$

Note next that *mutatis mutandis* (2.17), (2.18) and (2.23) all hold in the augmented model $\bar{\mathscr{P}}_N$ with score operator $\bar{\mathbb{A}}_\theta$ and tangent space H, and that the linear functional $\Psi(\cdot) = \langle \psi, \cdot \rangle_{\mathbb{H}}$ now verifies (2.16) as it is continuous on H for the $\|\bar{\mathbb{A}}_\theta [\cdot]\|_{L^2(\bar{P}_\theta)}$-norm, so that we can invoke the Riesz representation theorem to the effect that

$$\Psi(h) = \langle \bar{\mathbb{A}} \tilde{h}, \bar{\mathbb{A}} h \rangle_{L^2(\bar{P}_\theta)}, \quad h \in H, \text{ and some } \tilde{\psi}_\theta = \bar{\mathbb{A}} \tilde{h} \in \overline{(\bar{\mathscr{P}})_H}.$$

Thus the asymptotic minimax theorem in the augmented model gives

$$\liminf_{N \to \infty} \inf_{\bar{\psi}_N : (\mathbb{R} \times V \times \mathcal{X})^N \to \mathbb{R}} \sup_{h \in H, \|h\|_{\mathbb{H}} \le 1/\sqrt{N}} N E^N_{\theta+h} \big(\bar{\psi}_N - \Psi(\theta + h)\big)^2 \ge \bar{i}^{-1}_{\theta,H,\psi} = \sigma^2 \tag{2.25}$$

for estimators $\bar{\psi}$ based on the more informative data. The asymptotic local minimax risk in (2.24) exceeds the quantity in the last display, and letting $\sigma^2 \to \infty$ implies the result.

3. APPLICATION TO A DIVERGENCE FORM PDE

The results from the previous section describe how in a non-linear regression model (2.2) under Condition 1, the possibility of $\sqrt{N}$-consistent estimation of linear functionals $\Psi(\theta) = \langle \psi, \theta \rangle_{\mathbb{H}}$ essentially depends on whether ψ lies in the range of $\mathbb{I}_\theta^*$. A sufficient condition for this is that ψ lies in the range of the information operator $\mathbb{A}_\theta^* \mathbb{A}_\theta = \mathbb{I}_\theta^* \mathbb{I}_\theta$, and the results in **[20]** show that the lower bound in (2.18) can be attained by concrete estimators in this situation. The general theory was shown to apply to a class of PDEs of Schrödinger type **[20, 21]** and to non-linear X-ray transforms **[18, 20]**, with smooth test functions $\psi \in C^\infty$.

We now exhibit a PDE inverse problem where the range constraint from Theorem 2 fails, fundamentally limiting the possibility of efficient $\sqrt{N}$-consistent estimation of 'nice' linear functionals. In particular, we will show that, unlike for the Schrödinger type equations considered in **[20, 21]**, for this PDE the inverse Fisher information $\sigma_\theta^2(\psi)$ does not exist for a large class of functionals $\Psi(\theta) = \langle \theta, \psi \rangle_{L^2}$, including generic examples of *smooth non-negative* $\psi \in C^\infty$. This implies in particular the non-existence of a 'functional' Bernstein–von Mises phenomenon that would establish asymptotic normality of the posterior distribution of the process $\{\langle \theta, \psi \rangle_{L^2} : \psi \in C^\infty\}$ (comparable to those obtained in **[4, 5, 21]**).

3.1. Basic setting

Let $\mathcal{O} \subset \mathbb{R}^d$ be a bounded smooth domain with boundary $\partial\mathcal{O}$ and, for convenience, of unit volume $\lambda(\mathcal{O}) = 1$, where λ is Lebesgue measure. Denote by $C^\infty(\mathcal{O})$ the set of all smooth real-valued functions on $\mathcal{O}$ and by $C_0^\infty(\mathcal{O})$ the subspace of such functions of compact support in $\mathcal{O}$. Let $L^2 = L^2_\lambda(\mathcal{O})$ be the usual Hilbert space with inner product $\langle \cdot, \cdot \rangle_{L^2}$. The L^2_λ-Sobolev spaces $H^\beta = H^\beta(\mathcal{O})$ of order $\beta \in \mathbb{N}$ are also defined in the standard way, as are the spaces $C^\beta(\mathcal{O})$ that have all partial derivatives bounded and continuous up to order β.

For a *conductivity* $\theta \in C^\infty(\mathcal{O})$, *source* $f \in C^\infty(\mathcal{O})$ and *boundary temperatures* $g \in C^\infty(\partial\mathcal{O})$, consider solutions $u = u_\theta = u_{\theta,f,g}$ of the PDE

$$\begin{aligned} \nabla \cdot (\theta \nabla u) &= f \quad \text{in } \mathcal{O}, \\ u &= g \quad \text{on } \partial\mathcal{O}. \end{aligned} \tag{3.1}$$

Here $\nabla, \Delta, \nabla\cdot$ denote the gradient, Laplace and divergence operator, respectively. We ensure ellipticity by assuming $\theta \geq \theta_{\min} > 0$ throughout $\mathcal{O}$.

We write $\mathcal{L}_\theta = \nabla \cdot (\theta \nabla(\cdot))$ for the 'divergence form' operator featuring on the left-hand side in (3.1). A unique solution $u_\theta \in C^\infty(\mathcal{O})$ to (3.1) exists (e.g. Theorem 8.3 and Corollary 8.11 in [10]). The operator $\mathcal{L}_\theta$ has an inverse integral operator

$$V_\theta : L^2_\lambda(\mathcal{O}) \to H^2(\mathcal{O}) \cap \{h_{|\partial\mathcal{O}} = 0\} \tag{3.2}$$

for Dirichlet boundary conditions, that is, it satisfies $V_\theta[f] = 0$ at $\partial\mathcal{O}$ and $\mathcal{L}_\theta V_\theta[f] = f$ on $\mathcal{O}$ for all $f \in L^2_\lambda(\mathcal{O})$. Moreover, the operator V_θ is self-adjoint on $L^2_\lambda(\mathcal{O})$. One further shows that whenever $f \in H^2(\mathcal{O})$ satisfies $f_{|\partial\mathcal{O}} = 0$, then $V_\theta \mathcal{L}_\theta[f] = f$. These standard facts for elliptic PDEs can be proved, e.g. as in Section 5.1 in [29] or Chapter 2 in [17].

To define the 'forward map' $\mathcal{G}$ we consider a model Θ of conductivities arising as a H^β-neighbourhood of the standard Laplacian of radius $\eta > 0$, specifically

$$\Theta = \left\{ \theta \in C^\infty(\mathcal{O}), \inf_x \theta(x) > \frac{1}{2}, \theta_{|\partial\mathcal{O}} = 1 : \|\theta - 1\|_{H^\beta(\mathcal{O})} < \eta \right\}, \quad \beta > 1 + d. \tag{3.3}$$

The inverse problem is to recover θ from solutions

$$\mathcal{G} : \Theta \to L^2_\lambda(\mathcal{O}), \quad \mathcal{G}(\theta) \equiv u_\theta \tag{3.4}$$

of (3.1) where we emphasise that f, g, as well as $\theta_{|\partial\mathcal{O}}$, are assumed to be *known* (see also Remark 3). The particular numerical choices $1 = \theta_{|\partial\mathcal{O}}$ and $1/2 = \theta_{\min}$ are made for notational convenience. For independent $\varepsilon_i \sim^{iid} \mathcal{N}(0, 1)$, $X_i \sim^{iid} \lambda$, we then observe data

$$(Y_i, X_i)_{i=1}^N \in (\mathbb{R} \times \mathcal{O})^N \sim P_\theta^N, \quad Y_i = u_\theta(X_i) + \varepsilon_i, \tag{3.5}$$

from model (2.2). Note that unlike in statistical 'Calderón problems' [1], we measure u_θ throughout the entire domain $\mathcal{O}$. Before we take a closer look at the *local* information geometry of the map $\mathcal{G}$ arising from the PDE (3.1), let us first give conditions under which the problem of inferring θ from $(Y_i, X_i)_{i=1}^N$ in (3.5) has a consistent solution.

3.2. Global injectivity and model examples

Under suitable constellations of f, g in (3.1), the non-linear map $\theta \mapsto u_\theta$ can be injective, and 'stability' properties of $\mathscr{G}$ are well studied at least since [27], we refer to the recent contributions [2,13,24] and the many references therein. For instance, one can show:

Proposition 3. *Let $\theta_1, \theta_2 \in C^\infty(\mathcal{O})$ be conductivities such that $\|\theta_i\|_{C^1} \leq B$, $\theta_1 = \theta_2$ on $\partial\mathcal{O}$, and denote by u_{θ_i} the corresponding solutions to (3.1). Assume*

$$\inf_{x\in\mathcal{O}} \left[\Delta u_\theta(x) + \mu \left|\nabla u_\theta(x)\right|^2_{\mathbb{R}^d}\right] \geq c_0 > 0 \tag{3.6}$$

holds for $\theta = \theta_1$ and some $\mu > 0$. Then we have for some $C = C(B, \mu, c_0, \mathcal{O}) > 0$,

$$\|\theta_1 - \theta_2\|_{L^2} \leq C\|u_{\theta_1} - u_{\theta_2}\|_{H^2}. \tag{3.7}$$

Based on (3.7), one can show (see [13,24]) that we can recover θ in L^2-loss by some estimator $\hat{\theta} = \hat{\theta}((Y_i, X_i)_{i=1}^N)$ at a 'non-parametric rate' $\|\hat{\theta} - \theta\|_{L^2(\mathcal{O})} = O_{P_\theta^N}(N^{-\gamma})$ for some $0 < \gamma < 1/2$, uniformly in Θ. We wish to study here inference on linear functionals

$$\Psi(\theta) = \langle \psi, \theta\rangle_{L^2(\mathcal{O})}, \quad \psi \in C_0^\infty(\mathcal{O}).$$

As we can bound the 'plug-in' estimation error $|\langle \psi, \theta - \hat{\theta}\rangle_{L^2}|$ by $\|\hat{\theta} - \theta\|_{L^2}$, the convergence rate $N^{-\gamma}$ carries over to estimation of Ψ. Nevertheless, we will show that there are fundamental limitations for *efficient* inference on Ψ at the 'semi-parametric' rate ($\gamma = 1/2$). This will be illustrated with two model examples for which the 'injectivity' condition (3.6) can be checked.

Example 1 (No critical points). In (3.1), take

$$f = 2, \quad g = \frac{|\cdot|^2_{\mathbb{R}^d} - 1}{d}. \tag{3.8}$$

Then for the standard Laplacian $\theta = 1$, we have $u_1 = g$ on $\bar{\mathcal{O}}$, $\Delta u_1 = 2$, and hence $\nabla u_1 = 2x/d$, which satisfies $\inf_{x\in\mathcal{O}} |\nabla u_1(x)|_{\mathbb{R}^d} \geq c > 0$ for any domain $\mathcal{O} \subset \mathbb{R}^d$ *separated away from the origin*. This lower bound extends to

$$\inf_{\theta\in\Theta} \inf_{x\in\mathcal{O}} \left|\nabla u_\theta(x)\right|_{\mathbb{R}^d} \geq c_\nabla > 0 \tag{3.9}$$

for η small enough in (3.3), by perturbation: arguing as in (3.16) below and from standard elliptic regularity estimates (Lemma 23 in [24] and as in (3.15)), we have for $b > 1 + d/2$, $\beta > b + d/2$ (such that $H^\beta \subset C^b$),

$$\begin{aligned}\|u_\theta - u_1\|_{C^1} &\lesssim \left\|V_1\left[\nabla\cdot[(\theta - 1)\nabla u_\theta]\right]\right\|_{H^b} \lesssim \left\|(\theta - 1)\nabla u_\theta\right\|_{H^{b-1}} \\ &\lesssim \|\theta - 1\|_{H^{b-1}}\|u_\theta\|_{C^b} \leq \|\theta - 1\|_{H^\beta}\|u_\theta\|_{H^\beta} < C\eta.\end{aligned} \tag{3.10}$$

In view of $\sup_{\theta\in\Theta} \|\Delta u_\theta\|_\infty < \infty$ and (3.9), condition (3.6) is verified for μ large enough and all $\theta \in \Theta$.

The situation in Example 1 where the gradient ∇u_θ never vanishes is somewhat atypical, and one may expect u_θ to possess a *finite* number of isolated critical points x_0

(where $\nabla u_\theta(x_0)$ vanishes); see, e.g. [2] and references therein. The next example encompasses a prototypical such situation with an interior minimum. See also Remark 1 for the case of a saddle point. Further examples with more than one critical point are easily constructed, too.

Example 2 (Interior minimum). Consider the previous example where now $\mathcal{O}$ is the unit disk in $\mathbb{R}^2$ centred at the origin. In other words, in (3.1) we have $f = 2$ and $g_{|\partial\mathcal{O}} = 0$, corresponding to a classical Dirichlet problem with source f. In this case u_1 takes the same form as in the previous example but now has a gradient $\nabla u_1 = x$ that vanishes at the origin $0 \in \mathbb{R}^2$, corresponding to the unique minimum of u_1 on $\mathcal{O}$. The injectivity condition (3.6) is still satisfied for all $\theta \in \Theta$ simply since (3.1) implies

$$0 < 2 = \theta\Delta u_\theta + \nabla\theta \cdot \nabla u_\theta \quad \text{on } \mathcal{O},$$

so that either $\Delta u_\theta \geq 1/(2\|\theta\|_\infty)$ or $|\nabla u_\theta(x)|_{\mathbb{R}^d} \geq 1/(2\|\theta\|_{C^1})$ has to hold on $\mathcal{O}$. In this example, the constraints that η be small enough as well as that $\theta_1 = \theta_2$ on $\partial\mathcal{O}$ in Proposition 3 can in fact be removed, see Lemma 24 in [24].

3.3. The score operator and its adjoint

To connect to Section 2, let us regard Θ from (3.3) as a subset of the Hilbert space $\mathbb{H} = L^2_\lambda(\mathcal{O})$, and take $\mathscr{G}(\theta)$ from (3.4); hence we set $\mathcal{X} = \mathcal{O}$, $V = \mathbb{R}$, $\lambda = dx$ (Lebesgue measure).

As ‘tangent space’ $H \subset \mathbb{H}$, we take all smooth perturbations of θ of compact support,

$$H = C_0^\infty(\mathcal{O}), \tag{3.11}$$

so that the paths $\theta_{s,h} = \theta + sh$, $\theta \in \Theta$, $h \in H$, lie in Θ for all $s \in \mathbb{R}$ small enough. The closure $\bar{H}$ of H for $\|\cdot\|_{\mathbb{H}}$ equals $\bar{H} = \mathbb{H} = L^2_\lambda(\mathcal{O})$. We now check Condition 1, restricting to $d \leq 3$ to expedite the proof.

Theorem 4. *Assume $d \leq 3$. Let Θ be as in (3.3) and let the tangent space H be as in (3.11). The forward map $\theta \mapsto \mathscr{G}(\theta)$ from (3.4) satisfies Condition 1 for every $\theta \in \Theta$, with uniform bound $U_\mathscr{G} = U_\mathscr{G}(\|g\|_\infty, \|f\|_\infty)$ and with*

$$\mathbb{I}_\theta(h) \equiv -V_\theta\big[\nabla \cdot (h\nabla u_\theta)\big], \quad h \in H. \tag{3.12}$$

In particular, $\mathbb{I}_\theta$ extends to a bounded linear operator on $\mathbb{H}$.

Proof. We can represent the solutions u_θ of (3.1) by a Feynman–Kac-type formula as

$$u_\theta(x) = \mathbb{E}^x g(X_{\tau_\mathcal{O}}) - \mathbb{E}^x \int_0^{\tau_\mathcal{O}} f(X_s)ds, \quad x \in \mathcal{O}, \tag{3.13}$$

where $(X_s : s \geq 0)$ is a Markov diffusion process started at $x \in \mathcal{O}$ with infinitesimal generator $\mathcal{L}_\theta/2$, law $\mathbb{P}^x = \mathbb{P}^x_\theta$, and exit time $\tau_\mathcal{O}$ from $\mathcal{O}$, see Theorem 2.1 on p. 127 in [8]. As in the proof of Lemma 20 in [24], one bounds $\sup_{x\in\mathcal{O}} \mathbb{E}^x\tau_\mathcal{O}$ by a constant that depends only on $\mathcal{O}$, $\theta_{\min}$, and we conclude from the last display that therefore

$$\|u_\theta\|_\infty \le \|g\|_\infty + \|f\|_\infty \sup_{x\in\mathcal{O}} \mathbb{E}^x \tau_{\mathcal{O}} < \infty \tag{3.14}$$

so that the bound $U_{\mathscr{G}}$ for $\mathscr{G}$ required in Condition 1 follows.

We will repeatedly use the following elliptic regularity estimates:

$$\big\|V_\theta[h]\big\|_\infty \le c_0\big\|V_\theta[h]\big\|_{H^2} \le c_1\|h\|_{L^2}, \quad \|u_\theta\|_{H^2} \le c_2, \tag{3.15}$$

with constants $c_0 = c_0(\mathcal{O}), c_1 = c_1(\theta_{\min}, \mathcal{O}, \beta, \eta), c_2 = c_2(U_{\mathscr{G}}, \|f\|_{L^2}, \|g\|_{H^2}, \theta_{\min}, \mathcal{O}, \beta, \eta)$ that are *uniform* in $\theta \in \Theta$. The first inequality in (3.15) is just the Sobolev imbedding. The second follows from Lemma 21 in **[24]**, noting also that $\sup_{\theta\in\Theta}\|\theta\|_{C^1} \le C(\beta, \eta, \mathcal{O})$ by another Sobolev imbedding $H^\beta \subset C^1$. The final inequality in (3.15) follows from Theorem 8.12 in **[10]** and (3.14).

To verify (2.4), notice that the difference $u_{\theta+sh} - u_\theta$ solves (3.1) with $g = 0$ and appropriate right-hand side, specifically we can write

$$\mathscr{G}(\theta + sh) - \mathscr{G}(\theta) = -sV_\theta\big[\nabla\cdot(h\nabla u_{\theta+sh})\big], \quad h \in H, \tag{3.16}$$

for $|s|$ small enough. Then (2.4) follows from (3.15) since

$$\begin{aligned}\big\|V_\theta\big[\nabla\cdot(h\nabla u_{\theta+sh})\big]\big\|_\infty &\lesssim \big\|\nabla\cdot(h\nabla u_{\theta+sh})\big\|_{L^2} \lesssim \|h\nabla u_{\theta+sh}\|_{H^1}\\ &\lesssim \|h\|_{C^1}\sup_{\theta\in\Theta}\|u_\theta\|_{H^2} \le B < \infty.\end{aligned}$$

We will verify (2.3) by establishing a stronger '$\|\cdot\|_\infty$-norm' differentiability result: fix $\theta \in \Theta$ and any $h \in H$ such that $\theta + h \in \Theta$. Denote by $D\mathscr{G}_\theta[h]$ the solution $v = v_h$ of the PDE

$$\begin{aligned}\nabla\cdot(\theta\nabla v) &= -\nabla\cdot(h\nabla u_\theta) \quad \text{on } \mathcal{O},\\ v &= 0 \quad \text{on } \partial\mathcal{O}\end{aligned}$$

where u_θ is the given solution of the original PDE (3.1). Then the function $w_h = u_{\theta+h} - u_\theta - D\mathscr{G}_\theta[h]$ solves the PDE

$$\begin{aligned}\mathcal{L}_{\theta+h} w_h &= -\nabla\cdot\big(h\nabla V_\theta[\nabla\cdot(h\nabla u_\theta)]\big) \quad \text{on } \mathcal{O},\\ w_h &= 0 \quad \text{on } \partial\mathcal{O}.\end{aligned}$$

As a consequence, applying (3.15) and standard inequalities repeatedly, we have

$$\begin{aligned}\big\|u_{\theta+h} - u_\theta - D\mathscr{G}_\theta[h]\big\|_\infty &= \big\|V_{\theta+h}\big[\nabla\cdot\big(h\nabla V_\theta[\nabla\cdot(h\nabla u_\theta)]\big)\big]\big\|_\infty\\ &\lesssim \big\|\nabla\cdot\big(h\nabla V_\theta[\nabla\cdot(h\nabla u_\theta)]\big)\big\|_{L^2}\\ &\lesssim \|h\|_{C^1}\big\|V_\theta\big[\nabla\cdot(h\nabla u_\theta)\big]\big\|_{H^2}\\ &\lesssim \|h\|_{C^1}\big\|\nabla\cdot(h\nabla u_\theta)\big\|_{L^2}\\ &\lesssim \|h\|_{C^1}^2\|u_\theta\|_{H^2} = O\big(\|h\|_{C^1}^2\big).\end{aligned} \tag{3.17}$$

In particular $D\mathscr{G}_\theta[sh] = \mathbb{I}_\theta[sh]$ is the linearisation of the forward map $\theta \mapsto \mathscr{G}(\theta) = u_\theta$ along any path $\theta + sh$, $|s| > 0$, $h \in H$. Finally, by duality, self-adjointness of V_θ and the divergence theorem (Proposition 2.3 on p. 143 in [29]), we can bound for every $h \in H$,

$$\|\mathbb{I}_\theta h\|_{L^2} = \sup_{\|\phi\|_{L^2}\le 1} \left| \int_{\mathcal{O}} \phi V_\theta[\nabla \cdot (h \nabla u_\theta)] \right| = \sup_{\|\phi\|_{L^2}\le 1} \left| \int_{\mathcal{O}} \nabla V_\theta[\phi] \cdot h \nabla u_\theta \right|$$

$$\lesssim \sup_{\|\phi\|_{L^2}\le 1} \|V_\theta[\phi]\|_{H^1} \|h\|_{L^2} \|u_\theta\|_{C^1} \lesssim \|h\|_{L^2},$$

using also (3.15) and that $\|u_\theta\|_{C^1} < \infty$ (here for fixed θ) as u_θ is smooth. By continuity and since H is dense in $L^2_\lambda = \mathbb{H}$, we can extend $\mathbb{I}_\theta$ to a bounded linear operator on $\mathbb{H}$, completing the proof. ■

Theorem 1 gives the score operator $\mathbb{A}_\theta$ mapping H into $L^2(\mathbb{R} \times \mathcal{O}, P_\theta)$ of the form

$$\mathbb{A}_\theta[h](x, y) = \big(y - u_\theta(x)\big) \times \mathbb{I}_\theta(h)(x), \quad y \in \mathbb{R},\ x \in \mathcal{O}. \tag{3.18}$$

For the present tangent space H, we have $\bar{H} = \mathbb{H}$. To apply the general results from Section 2, we now calculate the adjoint $\mathbb{I}^*_\theta : L^2_\lambda(\mathcal{O}) \to \bar{H} = L^2_\lambda(\mathcal{O})$ of $\mathbb{I}_\theta : \bar{H} \to L^2(\mathcal{O})$.

Proposition 4. *The adjoint* $\mathbb{I}^*_\theta : L^2_\lambda(\mathcal{O}) \to L^2_\lambda(\mathcal{O})$ *of* $\mathbb{I}_\theta$ *is given by*

$$\mathbb{I}^*_\theta[g] = \nabla u_\theta \cdot \nabla V_\theta[g], \quad g \in L^2_\lambda(\mathcal{O}). \tag{3.19}$$

Proof. Since $\mathbb{I}_\theta$ from (3.12) defines a bounded linear operator on the Hilbert space $L^2_\lambda = \mathbb{H}$, a unique adjoint operator I^*_θ exists by the Riesz representation theorem. Let us first show that

$$\langle h, (I^*_\theta - \mathbb{I}^*_\theta) g \rangle_{L^2} = 0, \quad \forall h, g \in C^\infty_0(\mathcal{O}). \tag{3.20}$$

Indeed, since V_θ is self-adjoint for L^2_λ and satisfies $[V_\theta g]_{|\partial\mathcal{O}} = 0$, we can apply the divergence theorem (Proposition 2.3 on p. 143 in [29]) with vector field $X = h\nabla u_\theta$ to deduce

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

so that (3.20) follows. Since $C^\infty_0(\mathcal{O})$ is dense in $L^2_\lambda(\mathcal{O})$ and since I^*_θ, $\mathbb{I}^*_\theta$ are continuous on $L^2_\lambda(\mathcal{O})$ (by construction in the former case and by (3.15), $u_\theta \in C^\infty(\mathcal{O})$, in the latter case), the identity (3.20) extends to all $g \in L^2_\lambda(\mathcal{O})$ and hence $I^*_\theta = \mathbb{I}^*_\theta$, as desired. ■

Note further that for $\theta \in \Theta$ fixed, using (3.15), $u_\theta \in C^\infty$ and L^2-continuity of $\mathbb{I}_\theta$, we have $\|\mathbb{I}^*_\theta \mathbb{I}_\theta h\|_{H^1} \lesssim \|\mathbb{I}_\theta h\|_{L^2} \lesssim \|h\|_{L^2}$. The compactness of the embedding $H^1 \subset L^2$ now implies that the information operator $\mathbb{I}^*_\theta \mathbb{I}_\theta$ is a compact and self-adjoint operator on $L^2(\mathcal{O})$.

3.4. Injectivity of $\mathbb{I}_\theta$, $\mathbb{I}^*_\theta \mathbb{I}_\theta$

Following the developments in Section 2, our ultimate goal is to understand the range $R(\mathbb{I}^*_\theta)$ of the adjoint operator $\mathbb{I}^*_\theta$. A standard Hilbert space duality argument implies that

$$R(\mathbb{I}^*_\theta)^\perp = \ker(\mathbb{I}_\theta), \tag{3.21}$$

that is, the ortho-complement (in $\mathbb{H}$) of the range of $\mathbb{I}_\theta^*$ equals the kernel (null space) of $\mathbb{I}_\theta$ (in $\mathbb{H}$). Thus if ψ is in the kernel of $\mathbb{I}_\theta$ then it cannot lie in the range of the adjoint and the non-existence of the inverse Fisher information in Theorem 2 for such ψ can be attributed simply to the lack of injectivity of $\mathbb{I}_\theta$.

We first show that under the natural 'global identification' condition (3.6), the mapping $\mathbb{I}_\theta$ from (3.12) is injective on the tangent space H (and hence on our parameter space Θ). The proof (which is postponed to Section 4.1) also implies injectivity of the information operator $\mathbb{I}_\theta^*\mathbb{I}_\theta$ on H, and in fact gives an $H^2 - L^2$ Lipschitz stability estimate for $\mathbb{I}_\theta$.

Theorem 5. *In the setting of Theorem 4, suppose also that (3.6) holds true. Then for $\mathbb{I}_\theta$ from (3.12), every $\theta \in \Theta$ and some $c = c(\mu, c_0, \theta, \mathcal{O})$,*

$$\left\| \mathbb{I}_\theta[h] \right\|_{H^2} \geq c\|h\|_{L^2} \quad \forall h \in H. \tag{3.22}$$

In particular, $\mathbb{I}_\theta(h) = 0$ or $\mathbb{I}_\theta^*\mathbb{I}_\theta(h) = 0$ imply $h = 0$ for all $h \in H$.

Using (3.15), one shows further that the operator $\mathbb{I}_\theta$ is continuous from $H^1(\mathcal{O}) \to H^2(\mathcal{O})$ and, by taking limits in (3.22), Theorem 5 then extends to all $h \in H_0^1(\mathcal{O})$ obtained as the completion of H for the $H^1(\mathcal{O})$-Sobolev norm.

Of course, the kernel in (3.21) is calculated on the Hilbert space $\mathbb{H} = L^2(\mathcal{O})$, so the previous theorem does not characterise $R(\mathbb{I}_\theta^*)^\perp$, yet. Whether $\mathbb{I}_\theta$ is injective on all of $L^2(\mathcal{O})$ depends on finer details of the PDE (3.1). Let us illustrate this in the model examples from above.

3.4.1. Example 1 continued; on the kernel in $L^2(\mathcal{O})$

In our first example, $\mathbb{I}_\theta$ starts to have a kernel already when $h_{|\partial\mathcal{O}} \neq 0$. Indeed, from the proof of Theorem 5, a function $h \in C^\infty(\bar{\mathcal{O}})$ is in the kernel of $\mathbb{I}_\theta$ if and only if

$$T_\theta(h) = \nabla \cdot (h\nabla u_\theta) = \nabla h \cdot \nabla u_\theta + h\Delta u_\theta = 0. \tag{3.23}$$

Now fix any $\theta \in \Theta$ with u_θ satisfying (3.9). The integral curves $\gamma(t)$ in $\mathcal{O}$ associated to the smooth vector field $\nabla u_\theta \neq 0$ are given near $x \in \mathcal{O}$ as the unique solutions (e.g. [29, P. 9]) of the vector ODE

$$\frac{d\gamma}{dt} = \nabla u_\theta(\gamma), \quad \gamma(0) = x. \tag{3.24}$$

Since ∇u_θ does not vanish, we obtain through each $x \in \mathcal{O}$ a unique curve $(\gamma(t) : 0 \leq t \leq T_\gamma)$ originating and terminating at the boundary $\partial\mathcal{O}$, with finite 'travel time' $T_\gamma \leq T(\mathcal{O}, c_\nabla) < \infty$. Along this curve, (3.23) becomes the ODE

$$\frac{d}{dt}h(\gamma(t)) + h(\gamma(t))\Delta u_\theta(\gamma(t)) = 0, \quad 0 < t < T_\gamma.$$

Under the constraint $h|_{\partial\mathcal{O}} = 0$ for $h \in H$, the unique solution of this ODE is $h = 0$, which is in line with Theorem 5. But for other boundary values of h, non-zero solutions exist. One can characterise the elements $h \in C^\infty(\bar{\mathcal{O}})$ in the kernel of $\mathbb{I}_\theta$ as follows. Since the vector field ∇u_θ is non-trapping, there exists (see [7, THEOREM 6.4.1]) $r \in C^\infty(\bar{\mathcal{O}})$ such that $\nabla u_\theta \cdot \nabla r = \Delta u_\theta$. Thus

$$\nabla u_\theta \cdot \nabla(he^r) = e^r T_\theta(h)$$

and it follows that $T_\theta(h) = 0$ iff he^r is a first integral of ∇u_θ. Observe that the set of first integrals of ∇u_θ is rather large: using the flow of ∇u_θ, we can pick coordinates $(x_1, \dots, x_d)$ in $\mathcal{O}$ such that $t \mapsto (t + x_1, x_2, \dots, x_d)$ are the integral curves of ∇u_θ and thus any function that depends only on $x_2, \dots, x_d$ is a first integral.

3.4.2. Example 2 continued; injectivity on $L^2(\mathcal{O})$

We now show that in the context of Example 2, the injectivity part of Theorem 5 *does* extend to all of $L^2(\mathcal{O})$.

Proposition 5. *Let $\mathbb{I}_\theta$ be as in (3.12) where u_θ solves (3.1) with f, g as in (3.8) and $\mathcal{O}$ is the unit disk in $\mathbb{R}^2$ centred at $(0,0)$. Then for $\theta = 1$, the map $\mathbb{I}_1 : L^2(\mathcal{O}) \to L^2(\mathcal{O})$ is injective.*

Proof. Let us write $I = \mathbb{I}_1$ and suppose $I(f) = 0$ for $f \in L^2(\mathcal{O})$. Then for any $h \in C^\infty(\mathcal{O})$ we have by Proposition 4

$$0 = \langle If, h\rangle_{L^2(\mathcal{O})} = \langle f, I^* h\rangle_{L^2(\mathcal{O})} = \big\langle f, XV_1[h]\big\rangle_{L^2(\mathcal{O})} \tag{3.25}$$

with vector field $X = \nabla u_1 \cdot \nabla(\cdot) = x_1\partial x_1 + x_2 \partial x_2$, $(x_1, x_2) \in \mathcal{O}$. Choosing $h = \Delta g$ for any smooth g of compact support, we deduce that

$$\int_{\mathcal{O}} X(g) f \, d\lambda = 0, \quad \forall g \in C_0^\infty(\mathcal{O}), \tag{3.26}$$

and we now show that this implies $f = 0$. A somewhat informal dynamical argument would say that (3.26) 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 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 (3.26) as

$$\int_0^1 \left(r^2 \left(\int_0^{2\pi} f(r,\vartheta)\psi(\vartheta)\, d\vartheta \right) \partial_r \phi \right) dr = 0. \tag{3.27}$$

By Fubini's theorem, for each ψ 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 (3.27). 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 $L^2(\mathcal{O})$ and compute the pairing

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

unless $c_\psi = 0$. Thus $f = 0$. ∎

By perturbation (similar as in (3.10)) and the Morse lemma, we can show that $u_\theta, \theta \in \Theta$, has a gradient u_θ that vanishes only at a single point in a neighbourhood of 0, and so the proof of the previous theorem extends to any $\theta \in \Theta$.

3.5. The range of $\mathbb{I}_\theta^*$ and transport PDEs

From (3.21) we see $\overline{R(\mathbb{I}_\theta^*)} = \ker(\mathbb{I}_\theta)^\perp$, but in our infinite-dimensional setting care needs to be exercised as the last identity holds in the (complete) Hilbert space $\mathbb{H} = L^2(\mathcal{O})$ rather than in our tangent space H (on which the kernel of $\mathbb{I}_\theta$ is trivial). We will now show that the range $R(\mathbb{I}_\theta^*)$ remains strongly constrained. This is also true in Example 2 when $\ker(\mathbb{I}_\theta) = \{0\}$: the range may not be closed $\overline{R(\mathbb{I}_\theta^*)} \neq R(\mathbb{I}_\theta^*)$, and this 'gap' can be essential in the context of Theorems 2 and 3. To understand this, note that from Proposition 4 we have

$$R\big(\mathbb{I}_\theta^*\big) = \big\{\psi = \nabla u_\theta \cdot \nabla V_\theta[g], \text{ for some } g \in L^2_\lambda(\mathcal{O})\big\}. \tag{3.28}$$

The operator V_θ maps L^2_λ into $H^2_0 = \{y \in H^2 : y_{|\partial\mathcal{O}} = 0\}$ and hence if ψ is in the range of $\mathbb{I}_\theta^*$ then the equation

$$\begin{aligned} \nabla u_\theta \cdot \nabla y &= \psi \quad \text{on } \mathcal{O}, \\ y &= 0 \quad \text{on } \partial\mathcal{O} \end{aligned} \tag{3.29}$$

necessarily has a solution $y = y_\psi \in H^2_0$. The existence of solutions to the transport PDE (3.29) depends crucially on the compatibility of ψ with geometric properties of the vector field ∇u_θ, which in turn is determined by the geometry of the forward map $\mathscr{G}$ (via f, g, θ) in the base PDE (3.1). We now illustrate this in our two model Examples 1 and 2.

3.5.1. Example 1 continued; range constraint

Applying the chain rule to $y \in H^2(\mathcal{O})$ and using (3.24), we see

$$\frac{d}{dt} y\big(\gamma(t)\big) = \frac{d\gamma(t)}{dt} \cdot \nabla y\big(\gamma(t)\big) = (\nabla u_\theta \cdot \nabla y)\big(\gamma(t)\big), \quad 0 < t < T_\gamma.$$

Hence along any integral curve γ of the vector field ∇u_θ, the PDE (3.29) reduces to the ODE

$$\frac{dy}{dt} = \psi. \tag{3.30}$$

Now suppose $\psi \in R(\mathbb{I}_\theta^*)$. Then a solution $y \in H^2_0$ to (3.29) satisfying $y_{|\partial\mathcal{O}} = 0$ must exist. Such y then also solves the ODE (3.30) along each curve γ, with initial and terminal values $y(0) = y(T_\gamma) = 0$. By the fundamental theorem of calculus (and uniqueness of solutions), this forces

$$\int_0^{T_\gamma} \psi\big(\gamma(t)\big)\, dt = 0 \tag{3.31}$$

to vanish. In other words, ψ permits a solution y to (3.29) only if ψ integrates to zero along each integral curve (orbit) induced by the vector field ∇u_θ. Now consider any smooth (non-zero) *nonnegative* ψ in the tangent space $H = C_0^\infty(\mathcal{O})$, and take $x \in \mathcal{O}$ such that $\psi \geq c > 0$ near x. For γ the integral curve passing through x, we then cannot have (3.31) as the integrand never takes negative values while it is positive and continuous near x. Conclude by way of contradiction that $\psi \notin R(\mathbb{I}_\theta^*)$. Applying Theorems 2 and 3, we have proved:

Theorem 6. *Consider estimation of the functional $\Psi(\theta) = \langle \theta, \psi \rangle_{L^2(\mathcal{O})}$ from data $(Y_i, X_i)_{i=1}^N$ drawn i.i.d. from P_θ^N in the model (3.5) where f, g in (3.1) are chosen as in (3.8), the domain $\mathcal{O}$ is separated away from the origin, and Θ is as in (3.3) with η small enough and $\beta > 1 + d$,*

$d \le 3$. *Suppose* $0 \ne \psi \in C_0^\infty(\mathcal{O})$ *satisfies* $\psi \ge 0$ *on* $\mathcal{O}$. *Then for every* $\theta \in \Theta$, *the efficient Fisher information for estimating* $\Psi(\theta)$ *satisfies*

$$\inf_{h \in H, \langle h, \psi\rangle_{L^2} \ne 0} \frac{\|\mathbb{I}_\theta h\|^2_{L^2_\lambda}}{\langle \psi, h\rangle^2_{L^2_\lambda}} = 0. \tag{3.32}$$

In particular, for any $\theta \in \Theta$,

$$\liminf_{N\to\infty} \inf_{\tilde\psi_N:(\mathbb{R}\times\mathcal{O})^N\to\mathbb{R}} \sup_{\theta'\in\Theta, \|\theta'-\theta\|_{\mathbb{H}} \le 1/\sqrt{N}} N E^N_{\theta'}\left(\tilde\psi_N - \Psi(\theta')\right)^2 = \infty. \tag{3.33}$$

Let us notice that one can further show that (3.31) is also a *sufficient* condition for ψ to lie in the range of $\mathbb{I}^*_\theta$ (provided ψ is smooth and with compact support in $\mathcal{O}$). As this condition strongly depends on θ via the vector field ∇u_θ, it seems difficult to describe any choices of ψ that lie in $\bigcap_{\theta\in\Theta} R(\mathbb{I}^*_\theta)$.

3.5.2. Example 2 continued; range constraint

We showed in the setting of Example 2 that $\mathbb{I}_\theta$ is injective on all of $L^2(\mathcal{O})$, and hence any $\psi \in L^2(\mathcal{O})$ lies in *closure* of the range of $\mathbb{I}^*_\theta$. Nevertheless, there are many relevant ψ's that are not contained in $R(\mathbb{I}^*_\theta)$. In Example 2, the gradient of u_θ vanishes and the integral curves γ associated to $\nabla u_\theta = (x_1, x_2)$ emanate along straight lines from $(0,0)$ towards boundary points $(z_1, z_2) \in \partial\mathcal{O}$ where $y((z_1, z_2)) = 0$. If we parameterise them as $\{(z_1 e^t, z_2 e^t) : -\infty < t \le 0\}$, then as after (3.30) we see that if a solution $y \in H^2_0$ to (3.29) exists then ψ must necessarily satisfy

$$\int_{-\infty}^0 \psi\left(z_1 e^t, z_2 e^t\right) dt = 0 - y(0) = \text{const.} \quad \forall (z_1, z_2) \in \partial\mathcal{O}. \tag{3.34}$$

This again cannot happen, for example, for any non-negative non-zero $\psi \in H$ that vanishes along a given curve γ (for instance if it is zero in any given quadrant of $\mathcal{O}$), as this forces const $= 0$. Theorems 2 and 3 again yield the following for Example 2:

Theorem 7. *Consider the setting of Theorem 6 but where now* $\mathcal{O}$ *is the unit disk centred at* $(0,0)$, *and where* $0 \le \psi \in C_0^\infty(\mathcal{O})$, $\psi \ne 0$, *vanishes along some straight ray from* $(0,0)$ *to the boundary* $\partial\mathcal{O}$. *Then* (3.32) *and* (3.33) *hold at* $\theta = 1$.

Arguing as after Proposition 5, the result can be extended to any $\theta \in \Theta$ by an application of the Morse lemma.

3.6. Concluding remarks

Remark 1 (Interior saddle points of u_θ). To complement Examples 1, 2, suppose we take $\theta = 1$, $f = 0$ in (3.1) so that $u = u_1 = x_1^2 - x_2^2$ if $g = u_{\partial\mathcal{O}}$ (and $\mathcal{O}$ is the unit disk, say). Then $\nabla u = 2(x_1, -x_2)$ and the critical point is a saddle point. In this case we can find integral curves γ_x running through x away from $(0,0)$ between boundary points in finite time. Then is ψ is nonnegative and supported near x it cannot integrate to zero along γ_x. An analogue of Theorem 6 then follows for this constellation of parameters in (3.1), too. Note that in this example, the kernel of $\mathbb{I}_\theta$ contains at least all constants.

Remark 2 (Local curvature of $\mathscr{G}$). The quantitative nature of (3.22) in Theorem 5 is compatible with 'gradient stability conditions' employed in [3,25] to establish polynomial time posterior computation time bounds for gradient based Langevin MCMC schemes. Specifically, arguing as in Lemma 4.7 in [25], for a neighbourhood $\mathcal{B}$ of θ_0 one can deduce local average 'curvature'

$$\inf_{\theta \in \mathcal{B}} \lambda_{\min} E_{\theta_0}\big[-\nabla^2 \ell(\theta)\big] \geq c_2 D^{-4/d},$$

of the average-log-likelihood function ℓ when the model Θ is discretised in the eigenbasis $E_D \equiv (e_n : n \leq D) \subset H$ arising from the Dirichlet Laplacian. In this sense (using also the results from [13]) one can expect a Bayesian inference method based on data (2.2) and Gaussian process priors to be consistent and computable even in high-dimensional settings. This shows that such local curvature results are not sufficient to establish (and hence distinct from) Gaussian 'Bernstein–von Mises-type' approximations.

Remark 3 (Boundary constraints on θ). As the main flavour of our results is 'negative', the assumption of knowledge of the boundary values of θ in (3.3) strengthens our conclusions – it is also natural as the regression function $u = g$ is already assumed to be known at $\partial\mathcal{O}$. In the definition of the parameter space Θ, we could further have assumed that all outward normal derivatives up to order $\beta - 1$ of θ vanish at $\partial\mathcal{O}$. This would be in line with the parameter spaces from [13,24]. All results in this section remain valid because our choice of tangent space H in (3.11) is compatible with this more constrained parameter space.

Remark 4 (Ellipticity). The Bernstein–von Mises theorems from [18,20,21] exploit *ellipticity* of the information operator $\mathbb{I}_\theta^* \mathbb{I}_\theta$ in their settings, allowing one to solve for y in the equation $\mathbb{I}_\theta^* \mathbb{I}_\theta y = \psi$ so that $R(\mathbb{I}_\theta^*)$ contains at least all smooth compactly supported ψ (and this is so for *any* parameter $\theta \in \Theta$). In contrast, in the present inverse problem arising from (3.1), the information operator does not have this property and solutions y to the critical equation $\mathbb{I}_\theta^* y = \psi$ exist only under stringent geometric conditions on ψ. Moreover, these conditions exhibit a delicate dependence on θ, further constraining the set $\bigcap_{\theta \in \Theta} R(\mathbb{I}_\theta^*)$ relevant for purposes of statistical inference.

4. APPENDIX

For convenience of the reader we include here a few more proofs of some results of this article.

4.1. Proofs of Theorem 5 and Proposition 3

Define the operator

$$T_\theta(h) = \nabla \cdot (h \nabla u_\theta), \quad h \in H,$$

so that (3.12) becomes $\mathbb{I}_\theta = V_\theta \circ T_\theta$. The map $u \mapsto (\mathcal{L}_\theta u, u_{|\partial\mathcal{O}})$ is a topological isomorphism between $H^2(\mathcal{O})$ and $L^2(\mathcal{O}) \times H^{3/2}(\partial\mathcal{O})$ (see [17], Theorem II.5.4), and hence with $u = V_\theta[w]$ we deduce $\|V_\theta[w]\|_{H^2} \gtrsim \|w\|_{L^2}$ for all $w \in C^\infty(\mathcal{O})$. As a consequence, using also Lemma 1,

$$\|\mathbb{I}_\theta[h]\|_{H^2} \gtrsim \|T_\theta(h)\|_{L^2} \gtrsim \|h\|_{L^2}, \quad h \in H,$$

which proves the inequality in Theorem 5. Next, as $\mathbb{I}_\theta$ is linear, we see that whenever $\mathbb{I}_\theta[h_1] = \mathbb{I}_\theta[h_2]$ for $h_1, h_2 \in H$ we have $\mathbb{I}_\theta[h_1 - h_2] = 0$, and so by the preceding inequality $h = h_1 - h_2 = 0$ in L^2, too. Likewise, if $h_1, h_2 \in H$ are such that $\mathbb{I}_\theta^*\mathbb{I}_\theta h_1 = \mathbb{I}_\theta^*\mathbb{I}_\theta h_2$, then $0 = \langle \mathbb{I}_\theta^*\mathbb{I}_\theta(h_1 - h_2), h_1 - h_2\rangle_{L^2_\lambda} = \|\mathbb{I}_\theta(h_1 - h_2)\|^2_{L^2_\lambda}$ so $\mathbb{I}_\theta h_1 = \mathbb{I}_\theta h_2$ and thus by what precedes $h_1 = h_2$.

Lemma 1. *We have* $\|T_\theta(h)\|_{L^2} = \|\nabla\cdot(h\nabla u_\theta)\|_{L^2} \geq c\|h\|_{L^2}$ *for all* $h \in H$ *and some constant* $c = c(\mu, B, c_0) > 0$, *where* $B \geq \|u_\theta\|_\infty$.

Proof. Applying the Gauss–Green theorem to any $v \in C^1(\mathcal{O})$ vanishing at $\partial\mathcal{O}$ gives

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

For $v = e^{-\mu u_\theta}h$, $h \in H$, with $\mu > 0$ to be chosen, we thus have

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

so that by the Cauchy–Schwarz inequality

$$\begin{aligned}\left|\int_\mathcal{O}\left(\frac{1}{2}\Delta u_\theta + \mu\|\nabla u_\theta\|^2\right)v^2\right| &= \left|\left\langle\left(\Delta u_\theta + \mu\|\nabla u_\theta\|^2\right), v^2\right\rangle_{L^2} + \frac{1}{2}\langle\nabla u_\theta, \nabla(v^2)\rangle_{L^2}\right| \\ &= \left|\left\langle h\Delta u_\theta + \nabla h\cdot\nabla u_\theta, he^{-2\mu u_\theta}\right\rangle_{L^2}\right| \\ &\leq \bar{\mu}\left\|\nabla\cdot(h\nabla u_\theta)\right\|_{L^2}\|h\|_{L^2} \end{aligned} \tag{4.1}$$

for $\bar{\mu} = \exp(2\mu\|u_\theta\|_\infty)$. We next lower bound the multipliers of v^2 in the left-hand side of (4.1). By (3.6),

$$\left|\int_\mathcal{O}\left(\frac{1}{2}\Delta u_\theta + \mu\|\nabla u_\theta\|^2\right)v^2\right| \geq c_0\int_\mathcal{O} v^2$$

and, combining this with (4.1), we deduce

$$\left\|\nabla\cdot(h\nabla u_\theta)\right\|_{L^2}\|h\|_{L^2} \geq c'\|v\|^2_{L^2(\mathcal{O})} \gtrsim \|h\|^2_{L^2}, \; h \in H,$$

which is the desired estimate. ■

The last lemma also immediately implies Proposition 3. Let us write $h = \theta_1 - \theta_2$ which defines an element of H. Then by (3.1) we have $\nabla\cdot(h\nabla u_{\theta_1}) = \nabla\cdot(\theta_2\nabla(u_{\theta_2} - u_{\theta_1}))$ and hence $\|\nabla\cdot(h\nabla u_{\theta_1})\|_{L^2} \lesssim \|u_{\theta_2} - u_{\theta_1}\|_{H^2}$. By Lemma 1 the left-hand side is lower bounded by a constant multiple of $\|h\|_{L^2} = \|\theta_1 - \theta_2\|_{L^2}$, so that the result follows.

4.2. Proof of Theorem 2 for $\mathbb{I}_\theta^*\mathbb{I}_\theta$ compact

Let us assume $\bar{H} = \mathbb{H}$ without loss of generality, write $I \equiv \mathbb{I}_\theta$, $L^2 = L^2_\lambda(\mathcal{X})$ in this proof, and let $\ker(I^*I) = \{h \in \mathbb{H} : I^*Ih = 0\}$. If I^*I is a compact operator on $\mathbb{H}$ then by the spectral theorem for self-adjoint operators, there exists an orthonormal system of $\mathbb{H}$

of eigenvectors $\{e_k : k \in \mathbb{N}\}$ spanning $\mathbb{H} \ominus \ker(I^*I)$ corresponding to eigenvalues $\lambda_k > 0$ so that

$$I^*Ie_k = \lambda_k e_k, \quad \text{and} \quad I^*Ih = \sum_k \lambda_k \langle h, e_k\rangle_{\mathbb{H}} e_k, \quad h \in \mathbb{H}.$$

We can then define the usual square-root operator $(I^*I)^{1/2}$ by

$$\big(I^*I\big)^{1/2}h = \sum_k \lambda_k^{1/2}\langle h, e_k\rangle_{\mathbb{H}} e_k, \quad h \in \mathbb{H}. \tag{4.2}$$

If we denote by P_0 the $\mathbb{H}$-projection onto $\ker(I^*I)$, then the range of $(I^*I)^{1/2}$ equals

$$R\big(\big(I^*I\big)^{1/2}\big) = \Big\{g \in \mathbb{H} : P_0(g) = 0, \sum_k \lambda_k^{-1}\langle e_k, g\rangle_{\mathbb{H}}^2 < \infty\Big\}. \tag{4.3}$$

Indeed, using standard Hilbert space arguments, (a) since $P_0(e_k) = 0$ for all k, for any $h \in \mathbb{H}$ 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 $\sum_k \lambda_k^{-1}\langle e_k, g\rangle_{\mathbb{H}}^2 < \infty$ then $h = \sum_k \lambda_k^{-1/2}\langle e_k, g\rangle e_k$ belongs to $\mathbb{H}$ and $(I^*I)^{1/2}h = g$.

Next, Lemma A.3 in [30] implies that $R(I^*) = R((I^*I)^{1/2})$. Now suppose $\psi \in \mathbb{H}$ is such that $\psi \notin R(I^*)$ and hence $\psi \notin R((I^*I)^{1/2})$. Then from (4.3), either $P_0(\psi) \neq 0$ or $\sum_k \lambda_k^{-1}\langle e_k, \psi\rangle_{\mathbb{H}}^2 = \infty$ (or both). In the first case, let $\bar{h} = P_0(\psi)$, so

$$\|I\bar{h}\|_{L^2} = \big\|I\big(P_0(\psi)\big)\big\|_{L^2} = \big\langle I^*I\big(P_0(\psi)\big), P_0(\psi)\big\rangle_{\mathbb{H}} = 0,$$

but $\langle \psi, \bar{h}\rangle_{\mathbb{H}} = \|P_0\psi\|_{\mathbb{H}}^2 = \delta$ for some $\delta > 0$. Since H is dense in $\mathbb{H}$, for any $\epsilon, 0 < \epsilon < \min(\delta/(2\|\psi\|_{\mathbb{H}}), \delta^2/4)$, we can find $h \in H$ such that $\|h - \bar{h}\|_{\mathbb{H}} < \epsilon$ and by continuity also $\|I(h - \bar{h})\|_{L^2} < \epsilon$. Then

$$\sqrt{i_{\theta,h,\psi}} = \frac{\|Ih\|_{L^2}}{|\langle \psi, h\rangle_{\mathbb{H}}|} \leq 2\frac{\epsilon}{\delta} \leq \sqrt{\epsilon}.$$

Using also (2.12), we conclude that $i_{\theta,H,\psi} < \epsilon$ in (2.21), so that the result follows since ϵ was arbitrary. In the second case we have $\sum_k \lambda_k^{-1}\langle e_k, \psi\rangle_{\mathbb{H}}^2 = \infty$ and define

$$\psi_N = \sum_{k \leq N} \lambda_k^{-1} e_k \langle e_k, \psi\rangle_{\mathbb{H}}, \quad N \in \mathbb{N},$$

which defines an element of $\mathbb{H}$. By density we can choose $h_N \in H$ such that $\|h_N - \psi_N\|_{\mathbb{H}} < 1/\|\psi\|_{\mathbb{H}}$, as well as $\|I(h_N - \psi_N)\|_{L^2} < 1$, for every N fixed. Next observe that

$$\langle \psi, \psi_N\rangle_{\mathbb{H}} = \sum_{k \leq N} \lambda_k^{-1}\langle e_k, \psi\rangle_{\mathbb{H}}^2 \equiv M_N,$$

$$\big\|I(\psi_N)\big\|_{L^2}^2 = \big\langle I^*I(\psi_N), \psi_N\big\rangle_{\mathbb{H}} = \sum_{k \leq N} \lambda_k^{-1}\langle e_k, \psi\rangle_{\mathbb{H}}^2 = M_N,$$

and that $M_N \to \infty$ as $N \to \infty$. Then by our choice of $h_N \in H$ and if $M_N \geq 2$, we have by the triangle inequality,

$$\big|\langle \psi, h_N\rangle_{\mathbb{H}}\big| \geq \big|\langle \psi, \psi_N\rangle_{\mathbb{H}}\big| - \big|\langle \psi, \psi_N - h_N\rangle_{\mathbb{H}}\big| \geq M_N - 1 \geq M_N/2,$$

$$\big\|I(h_N)\big\|_{L^2} \leq \big\|I(\psi_N)\big\|_{L^2} + \big\|I(h_N - \psi_N)\big\|_{L^2} \leq \sqrt{M}_N + 1 \leq 2\sqrt{M}_N.$$

From this and (2.12) we conclude that the inverse of (2.21) satisfies

$$i_{\theta,H,\psi}^{-1} \geq \frac{\langle \psi, h_N \rangle_{\mathbb{H}}^2}{\|I h_N\|_{L^2}^2} \geq \frac{1}{16} \frac{M_N^2}{M_N} \geq M_N/16.$$

As N was arbitrary and $M_N \to_{N\to\infty} \infty$, we must have $i_{\theta,H,\psi} = 0$, as desired.

ACKNOWLEDGMENTS

We are grateful to Jan Bohr and Lauri Oksanen for helpful remarks and discussions.

REFERENCES

[1] K. Abraham and R. Nickl, On statistical Caldéron problems. *Math. Statist. Learn.* **2** (2019), no. 2, 165–216.

[2] G. S. Alberti, G. Bal, and M. Di Cristo, Critical points for elliptic equations with prescribed boundary conditions. *Arch. Ration. Mech. Anal.* **226** (2017), no. 1, 117–141.

[3] J. Bohr and R. Nickl, On log-concave approximations of high-dimensional posterior measures and stability properties in non-linear inverse problems. 2021, arXiv:2105.07835.

[4] I. Castillo and R. Nickl, Nonparametric Bernstein–von Mises theorems in Gaussian white noise. *Ann. Statist.* **41** (2013), no. 4, 1999–2028.

[5] I. Castillo and R. Nickl, On the Bernstein–von Mises phenomenon for nonparametric Bayes procedures. *Ann. Statist.* **42** (2014), no. 5, 1941–1969.

[6] I. Castillo and J. Rousseau, A Bernstein–von Mises theorem for smooth functionals in semiparametric models. *Ann. Statist.* **43** (2015), no. 6, 2353–2383.

[7] J. J. Duistermaat and L. Hörmander, Fourier integral operators. II. *Acta Math.* **128** (1972), no. 3–4, 183–269.

[8] M. Freidlin, *Functional integration and partial differential equations*. Ann. of Math. Stud. 109, Princeton University Press, Princeton, NJ, 1985.

[9] C.-F. Gauß, *Theoria Motus Corporum Coelestium*. Perthes, Hamburg, 1809.

[10] D. Gilbarg and N. S. Trudinger, *Elliptic partial differential equations of second order*. Springer, Berlin–New York, 1998.

[11] E. Giné and R. Nickl, *Mathematical foundations of infinite-dimensional statistical models*. Cambridge University Press, New York, 2016.

[12] M. Giordano and H. Kekkonen, Bernstein–von Mises theorems and uncertainty quantification for linear inverse problems. *SIAM/ASA J. Uncertain. Quantificat.* **8** (2020).

[13] M. Giordano and R. Nickl, Consistency of Bayesian inference with Gaussian process priors in an elliptic inverse problem. *Inverse Probl.* (2020).

[14] M. Giordano and K. Ray, Nonparametric Bayesian inference for reversible multidimensional diffusions. 2018, arXiv:1802.05635.

[15] 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.

[16] H. Kekkonen, Consistency of Bayesian inference with Gaussian process priors for a parabolic inverse problem. 2021, arXiv:2103.13213.

[17] J.-L. Lions and E. Magenes, *Non-homogeneous boundary value problems and applications. Vol. I*. Springer, New York–Heidelberg, 1972.

[18] F. Monard, R. Nickl, and G. P. Paternain, Efficient nonparametric Bayesian inference for X-ray transforms. *Ann. Statist.* **47** (2019), no. 2, 1113–1147.

[19] F. Monard, R. Nickl, and G. P. Paternain, Consistent inversion of noisy non-abelian X-ray transforms. *Comm. Pure Appl. Math.* **74** (2021), 1045–1099.

[20] F. Monard, R. Nickl, and G. P. Paternain, Statistical guarantees for Bayesian uncertainty quantification in non-linear inverse problems with Gaussian process priors. *Ann. Statist.* **49** (2021), 3255–3298.

[21] R. Nickl, Bernstein–von Mises theorems for statistical inverse problems I: Schrödinger equation. *J. Eur. Math. Soc. (JEMS)* **22** (2020), 2697–2750.

[22] R. Nickl and K. Ray, Nonparametric statistical inference for drift vector fields of multi-dimensional diffusions. *Ann. Statist.* **48** (2020), no. 3, 1383–1408.

[23] R. Nickl and J. Söhl, Nonparametric Bayesian posterior contraction rates for discretely observed scalar diffusions. *Ann. Statist.* **45** (2017), no. 4, 1664–1693.

[24] R. Nickl, S. van de Geer, and S. Wang, Convergence rates for penalised least squares estimators in PDE-constrained regression problems. *SIAM/ASA J. Uncertain. Quantificat.* **8** (2020).

[25] R. Nickl and S. Wang, On polynomial-time computation of high-dimensional posterior measures by Langevin-type algorithms. 2020, arXiv:2009.05298.

[26] K. Ray, Adaptive Bernstein–von Mises theorems in Gaussian white noise. *Ann. Statist.* **45** (2017), no. 6, 2511–2536.

[27] G. R. Richter, An inverse problem for the steady state diffusion equation. *SIAM J. Appl. Math.* **41** (1981), no. 2, 210–221.

[28] A. M. Stuart, Inverse problems: a Bayesian perspective. *Acta Numer.* **19** (2010), 451–559.

[29] M. E. Taylor, *Partial differential equations I. Basic theory*. 2nd edn., Appl. Math. Sci. 115, Springer, New York, 2011.

[30] A. W. van der Vaart, On differentiable functionals. *Ann. Statist.* **19** (1991), no. 1, 178–204.

[31] A. W. van der Vaart, *Asymptotic statistics*. Cambridge Univ. Press, 1998.

RICHARD NICKL
University of Cambridge, Faculty of Mathematics, Cambridge CB3 0WA, UK,
nickl@maths.cam.ac.uk

GABRIEL P. PATERNAIN
University of Cambridge, Faculty of Mathematics, Cambridge CB3 0WA, UK,
g.p.paternain@dpmms.cam.ac.uk