Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/135643>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Random forward models and log-likelihoods in Bayesian inverse problems

H. C. Lie^{*} T. J. Sullivan^{†‡} A. L. Teckentrup^{§¶}

October 1, 2018

Abstract: We consider the use of randomised forward models and log-likelihoods within the Bayesian approach to inverse problems. Such random approximations to the exact forward model or log-likelihood arise naturally when a computationally expensive model is approximated using a cheaper stochastic surrogate, as in Gaussian process emulation (kriging), or in the field of probabilistic numerical methods. We show that the Hellinger distance between the exact and approximate Bayesian posteriors is bounded by moments of the difference between the true and approximate log-likelihoods. Example applications of these stability results are given for randomised misfit models in large data applications and the probabilistic solution of ordinary differential equations.

Keywords: Bayesian inverse problem, random likelihood, surrogate model, posterior consistency, uncertainty quantification, randomised misfit, probabilistic numerics.

2010 Mathematics Subject Classification: 62F15, 62G08, 65C99, 65D05, 65D30, 65J22, 68W20.

1 Introduction

Inverse problems are ubiquitous in the applied sciences and in recent years renewed attention has been paid to their mathematical and statistical foundations (Evans and Stark, 2002; Kaipio and Somersalo, 2005; Stuart, 2010). Questions of well-posedness — i.e. the existence, uniqueness, and stability of solutions — have been of particular interest for infinite-dimensional/non-parametric inverse problems because of the need to ensure stable and discretisation-independent inferences (Lassas and Siltanen, 2004) and develop algorithms that scale well with respect to high discretisation dimension (Cotter et al., 2013).

This paper considers the stability of the posterior distribution in a Bayesian inverse problem (BIP) when an accurate but computationally intractable forward model or likelihood is replaced by a random surrogate or emulator. Such stochastic surrogates arise often in practice. For example, an expensive forward model such as the solution of a PDE may be replaced by a kriging/Gaussian process (GP) model (Stuart and Teckentrup, 2017). In the realm of “big data” a residual vector of prohibitively high dimension may be randomly subsampled or orthogonally projected onto a randomly-chosen low-dimensional subspace (Le et al., 2017; Nemirovski et al., 2008). In the field of probabilistic numerical methods (Hennig et al., 2015), a deterministic dynamical system may be solved stochastically, with the stochasticity representing epistemic uncertainty about the behaviour of the system below the temporal or spatial grid scale (Conrad et al., 2016; Lie et al., 2017).

In each of the above-mentioned settings, the stochasticity in the forward model propagates to associated inverse problems, so that the Bayesian posterior becomes a *random measure*, μ_N^S , which we define

^{*}Institute of Mathematics, Freie Universität Berlin, Arnimallee 6, 14195 Berlin, Germany, hlie@math.fu-berlin.de

[†]Institute of Mathematics, Freie Universität Berlin, Arnimallee 6, 14195 Berlin, Germany, t.j.sullivan@fu-berlin.de

[‡]Zuse Institute Berlin, Takustraße 7, 14195 Berlin, Germany, sullivan@zib.de

[§]School of Mathematics, University of Edinburgh, UK, a.teckentrup@ed.ac.uk

[¶]The Alan Turing Institute, 96 Euston Road, London, NW1 2DB, UK

precisely in (3.1). Alternatively, one may choose to average over the randomness to obtain a *marginal posterior*, μ_N^M , which we define precisely in (3.2). It is natural to ask in which sense the approximate posterior (either the random or the marginal version) is close to the ideal posterior of interest, μ .

In earlier work, Stuart and Teckentrup (2017) examined the case in which the random surrogate was a GP. More precisely, the object subjected to GP emulation was either the forward model (i.e. the parameter-to-observation map) or the negative log-likelihood. The prior GP was assumed to be continuous, and was then conditioned upon finitely many observations (i.e. pointwise evaluations) of the parameter-to-observation map or negative log-likelihood as appropriate. That paper provided error bounds on the Hellinger distance between the BIP's exact posterior distribution and various approximations based on the GP emulator, namely approximations based on the mean of the predictive (i.e. conditioned) GP, as well as approximations based on the full GP emulator. Those results showed that the Hellinger distance between the exact BIP posterior and its approximations can be bounded by moments of the error in the emulator.

In this paper, we extend the analysis of Stuart and Teckentrup (2017) to consider more general (i.e. non-Gaussian) random approximations to forward models and log-likelihoods, and quantify the impact upon the posterior measure in a BIP. After establishing some notation in Section 2, we state the main approximation theorems in Section 3. Section 4 gives an application of the general theory to random misfit models, in which high-dimensional data are rendered tractable by projection into a randomly-chosen low-dimensional subspace. Section 5 gives an application to the stochastic numerical solution of deterministic dynamical systems, in which the stochasticity is a device used to represent the impact of numerical discretisation uncertainty. The proofs of all theorems are deferred to an appendix located after the bibliographic references.

2 Setup and notation

2.1 Spaces of probability measures

Throughout, $(\Omega, \mathcal{F}, \mathbb{P})$ is a fixed probability space that is rich enough to serve as a common domain for all random variables of interest.

The space of probability measures on the Borel σ -algebra of a topological space $\mathcal{U}$ will be denoted by $\mathcal{M}_1(\mathcal{U})$; in practice, $\mathcal{U}$ will be a separable Banach space.

When $\mu \in \mathcal{M}_1(\mathcal{U})$, integration of a measurable function (random variable) $f: \mathcal{U} \rightarrow \mathbb{R}$ will also be denoted by expectation, i.e. $\mathbb{E}_\mu[f] := \int_{\mathcal{U}} f(u) d\mu(u)$.

The space $\mathcal{M}_1(\mathcal{U})$ will be endowed with the Hellinger metric $d_H: \mathcal{M}_1(\mathcal{U})^2 \rightarrow \mathbb{R}_{\geq 0}$: for $\mu, \nu \in \mathcal{M}_1(\mathcal{U})$ that are both absolutely continuous with respect to a reference measure π ,

$$d_H(\mu, \nu)^2 := \frac{1}{2} \int_{\mathcal{U}} \left| \sqrt{\frac{d\mu}{d\pi}} - \sqrt{\frac{d\nu}{d\pi}} \right|^2 d\pi = 1 - \int_{\mathcal{U}} \sqrt{\frac{d\mu}{d\pi} \frac{d\nu}{d\pi}} d\pi = 1 - \mathbb{E}_\nu \left[\sqrt{\frac{d\mu}{d\nu}} \right]. \quad (2.1)$$

The Hellinger distance is in fact independent of the choice of reference measure π and defines a metric on $\mathcal{M}_1(\mathcal{U})$ (Bogachev, 2007, Lemma 4.7.35–36) with respect to which $\mathcal{M}_1(\mathcal{U})$ evidently has diameter at most 1. The Hellinger topology coincides with the total variation topology (Kraft, 1955) and is strictly weaker than the Kullback–Leibler (relative entropy) topology (Pinsker, 1964); all these topologies are strictly stronger than the topology of weak convergence of measures.

As used in Sections 3–5, the Hellinger metric is useful for uncertainty quantification when assessing the similarity of Bayesian posterior probability distributions, since expected values of square-integrable functions are Lipschitz continuous with respect to the Hellinger metric:

$$|\mathbb{E}_\mu[f] - \mathbb{E}_\nu[f]| \leq 2 \sqrt{\mathbb{E}_\mu[|f|^2] + \mathbb{E}_\nu[|f|^2]} d_H(\mu, \nu) \quad (2.2)$$

when $f \in L_\mu^2(\mathcal{U}) \cap L_\nu^2(\mathcal{U})$. In particular, for bounded f , $|\mathbb{E}_\mu[f] - \mathbb{E}_\nu[f]| \leq 2\sqrt{2}\|f\|_\infty d_H(\mu, \nu)$.

2.2 Bayesian inverse problems

By an *inverse problem* we mean the recovery of $u \in \mathcal{U}$ from an imperfect observation $y \in \mathcal{Y}$ of $G(u)$, for a known forward operator $G: \mathcal{U} \rightarrow \mathcal{Y}$. In practice, the operator G may arise as the composition

$G = O \circ S$ of the solution operator $S: \mathcal{U} \rightarrow \mathcal{V}$ of a system of ordinary or partial differential equations with an observation operator $O: \mathcal{V} \rightarrow \mathcal{Y}$, and it is typically the case that $\mathcal{Y} = \mathbb{R}^J$ for some $J \in \mathbb{N}$, whereas $\mathcal{U}$ and $\mathcal{V}$ can have infinite dimension. For simplicity, we assume an additive noise model

$$y = G(u) + \eta, \quad (2.3)$$

where the statistics but not the realisation of η are known. In the strict sense, this inverse problem is ill-posed in the sense that there may be no element $u \in \mathcal{U}$ for which $G(u) = y$, or there may be multiple such u that are highly sensitive to the observed data y .

The Bayesian perspective eases these problems by interpreting u , y , and η all as random variables or fields. Through knowledge of the distribution of η , (2.3) defines the conditional distribution of $y|u$. After positing a prior probability distribution $\mu_0 \in \mathcal{M}_1(\mathcal{U})$ for u , the Bayesian solution to the inverse problem is nothing other than the posterior distribution for the conditioned random variable $u|y$. This posterior measure, which we denote $\mu^y \in \mathcal{M}_1(\mathcal{U})$, is from the Bayesian point of view the proper synthesis of the prior information in μ_0 with the observed data y . The same posterior μ^y can also be arrived at via the minimisation of penalised Kullback–Leibler, χ^2 , or Dirichlet energies (Dupuis and Ellis, 1997; Jordan and Kinderlehrer, 1996; Ohta and Takatsu, 2011), where the penalisation again expresses compromise between fidelity to the prior and fidelity to the data.

The rigorous formulation of Bayes' formula for this context requires careful treatment and some further notation (Stuart, 2010). The pair (u, y) is assumed to be a well-defined random variable with values in $\mathcal{U} \times \mathcal{Y}$. The marginal distribution of u is the Bayesian prior $\mu_0 \in \mathcal{M}_1(\mathcal{U})$. The observational noise η is distributed according to $\mathbb{Q}_0 \in \mathcal{M}_1(\mathcal{Y})$, independently of u . The random variable $y|u$ is distributed according to $\mathbb{Q}_u$, the translate of $\mathbb{Q}_0$ by $G(u)$, which is assumed to be absolutely continuous with respect to $\mathbb{Q}_0$, with

$$\frac{d\mathbb{Q}_u}{d\mathbb{Q}_0}(y) \propto \exp(-\Phi(u; y)).$$

The function $\Phi: \mathcal{U} \times \mathcal{Y} \rightarrow \mathbb{R}$ is called the *negative log-likelihood* or simply *potential*. In the elementary setting of centred Gaussian noise, $\eta \sim \mathcal{N}(0, \Gamma)$ on $\mathcal{Y} = \mathbb{R}^J$, the potential is the non-negative quadratic misfit¹ $\Phi(u; y) = \frac{1}{2} \|\Gamma^{-1/2}(y - G(u))\|_{\mathcal{Y}}^2$. However, particularly for cases in which $\dim \mathcal{Y} = \infty$, it may be necessary to allow Φ to take negative values and even to be unbounded below (Stuart, 2010, Remark 3.8).

With this notation, Bayes' theorem is then as follows (Dashti and Stuart, 2016, Theorem 3.4):

Theorem 2.1 (Generalised Bayesian formula). *Suppose that $\Phi: \mathcal{U} \times \mathcal{Y} \rightarrow \mathbb{R}$ is $\mu_0 \otimes \mathbb{Q}_0$ -measurable and that*

$$Z(y) := \mathbb{E}_{\mu_0}[\exp(-\Phi(u; y))]$$

satisfies $0 < Z(y) < \infty$ for $\mathbb{Q}_0$ -almost all $y \in \mathcal{Y}$. Then, for such y , the conditional distribution μ^y of $u|y$ exists and is absolutely continuous with respect to μ_0 with density

$$\frac{d\mu^y}{d\mu_0}(u) = \frac{\exp(-\Phi(u; y))}{Z(y)}. \quad (2.4)$$

Note that, for (2.4) to make sense, it is essential to check that $0 < Z(y) < \infty$. Hereafter, to save space, we regard the data y as fixed, and hence write $\Phi(u)$ in place of $\Phi(u; y)$, Z in place of $Z(y)$, and μ in place of μ^y . In particular, we shall redefine the negative log-likelihood as a function $\Phi: \mathcal{U} \rightarrow \mathbb{R}$, instead of a function $\Phi: \mathcal{U} \times \mathcal{Y} \rightarrow \mathbb{R}$ as in Theorem 2.1 above.

From the perspective of numerical analysis, it is natural to ask about the well-posedness of the Bayesian posterior μ : is it stable when the prior μ_0 , the potential Φ , or the observed data y are slightly perturbed, e.g. due to discretisation, truncation, or other numerical errors? For example, what is the impact of using an approximate numerical forward operator G_N in place of G , and hence an approximate $\Phi_N: \mathcal{U} \rightarrow \mathbb{R}$ in place of Φ ? Here, we quantify stability in the Hellinger metric d_H from (2.1).

Stability of the posterior with respect to the observed data y and the log-likelihood Φ was established for Gaussian priors by Stuart (2010) and for more general priors by many later contributions (Dashti et al., 2012; Hosseini, 2017; Hosseini and Nigam, 2017; Sullivan, 2017). (We note in passing that the stability of BIPs with respect to perturbation of the prior is possible but much harder to establish, particularly

¹Hereafter, to reduce notational clutter, we write both $\|\cdot\|_{\mathcal{U}}$ and $\|\cdot\|_{\mathcal{Y}}$ as $\|\cdot\|$.

when the data y are highly informative and the normalisation constant $Z(y)$ is close to zero; see e.g. the “brittleness” phenomenon of (Owhadi and Scovel, 2017; Owhadi et al., 2015). Typical approximation theorems for the replacement of the potential Φ by a deterministic approximate potential Φ_N , leading to an approximate posterior μ_N , aim to transfer the convergence rate of the forward problem to the inverse problem, i.e. to prove an implication of the form

$$|\Phi(u) - \Phi_N(u)| \leq M(\|u\|)\psi(N) \implies d_H(\mu, \mu_N) \leq C\psi(N),$$

where $M: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is suitably well behaved, $\psi: \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ quantifies the convergence rate of the forward problem, and C is a constant. Following Stuart and Teckentrup (2017), the purpose of this article is to extend this paradigm and these approximation results to the case in which the approximation Φ_N is a *random* object.

3 Well-posed Bayesian inverse problems with random likelihoods

In many practical applications, the negative log-likelihood Φ is computationally too expensive or impossible to evaluate exactly; one therefore often uses an approximation Φ_N of Φ . This leads to an approximation μ_N of the exact posterior μ , and a key desideratum is convergence, in a suitable sense, of μ_N to μ as the approximation error $\Phi_N - \Phi$ in the potential tends to zero.

The focus of this work is on random approximations Φ_N . One particular example of such random approximations are the GP emulators analysed in Stuart and Teckentrup (2017); other examples include the randomised misfit models in Section 4 and the probabilistic numerical methods in Section 5. The present section extends the analysis of Stuart and Teckentrup (2017) from the case of GP approximations of forward models or log-likelihoods to more general non-Gaussian approximations. In doing so, more precise conditions are obtained for the exact Bayesian posterior to be well approximated by its random counterpart.

Let now $\Phi_N: \Omega \times \mathcal{U} \rightarrow \mathbb{R}$ be a measurable function that provides a random approximation to $\Phi: \mathcal{U} \rightarrow \mathbb{R}$, where we recall that we have fixed the data y . Let ν_N be a probability measure on Ω such that the distribution of the inputs of Φ_N is given by $\nu_N \otimes \mu_0$; we sometimes abuse notation and think of Φ_N itself as being ν_N -distributed. We assume throughout that the randomness in the approximation Φ_N of Φ is independent of the randomness in the parameters being inferred.

Replacing Φ by Φ_N in (2.4), we obtain the *sample approximation* μ_N^S , the random measure given by

$$\frac{d\mu_N^S}{d\mu_0}(\omega, u) := \frac{\exp(-\Phi_N(\omega, u))}{Z_N^S}, \quad (3.1)$$

$$Z_N^S(\omega) := \mathbb{E}_{\mu_0}[\exp(-\Phi_N(\omega, \cdot))] = \int_{\mathcal{U}} \exp(-\Phi_N(\omega, u')) d\mu_0(u').$$

(Henceforth, we will omit the ω argument for brevity.) Thus, the measure μ is approximated by the random measure $\mu_N^S: \Omega \rightarrow \mathcal{M}_1(\mathcal{U})$, and the normalisation constant $Z_N^S: \Omega \rightarrow \mathbb{R}$ is a random variable. A deterministic approximation of the posterior distribution μ can now be obtained either by fixing ω , i.e. by taking one particular realisation of the random posterior μ_N^S , or by taking the expected value of the random likelihood $\exp(-\Phi_N(u))$, i.e. by averaging over different realisations of μ_N^S . This yields the *marginal approximation* μ_N^M defined by

$$\frac{d\mu_N^M}{d\mu_0}(u) := \frac{\mathbb{E}_{\nu_N}[\exp(-\Phi_N(u))]}{\mathbb{E}_{\nu_N}[Z_N^S]}, \quad (3.2)$$

where $\mathbb{E}_{\nu_N}[Z_N^S] = \int_{\Omega} Z_N^S(\omega) d\nu_N(\omega)$. We note that an alternative averaged, deterministic approximation can be obtained by taking the expected value of the density $(Z_N^S)^{-1} e^{-\Phi_N(u)}$ in (3.1) as a whole, i.e. by taking the expected value of the ratio rather than the ratio of expected values. A result very similar to Theorem 3.1, with slightly modified assumptions, holds also in this case, with the proof following the same steps. However, the marginal approximation presented here appears more intuitive and more amenable to applications. Firstly, the marginal approximation provides a clear interpretation as the posterior distribution obtained by the approximation of the true data likelihood $\exp(-\Phi(u))$ by $\mathbb{E}_{\nu_N}[\exp(-\Phi_N(u))]$. Secondly, the marginal approximation is more amenable to sampling methods such as Markov chain Monte

Carlo, with clear connections to the pseudo-marginal approach (Andrieu and Roberts, 2009; Beaumont, 2003).

3.1 Random misfit models

This section considers the general setting in which the deterministic potential Φ is approximated by a random potential $\Phi_N \sim \nu_N$. Recall from (2.4) that Z is the normalisation constant of μ , and that for μ to be well-defined, we must have that $0 < Z < \infty$. The following two results, Theorems 3.1 and 3.2, extend Theorems 4.9 and 4.11 respectively of Stuart and Teckentrup (2017), in which the approximation is a GP model:

Theorem 3.1 (Deterministic convergence of the marginal posterior). *Suppose that there exist scalars $C_1, C_2, C_3 \geq 0$, independent of N , such that, for the Hölder-conjugate exponent pairs (p_1, p'_1) , (p_2, p'_2) , and (p_3, p'_3) , we have*

- (a) $\min \left\{ \left\| \mathbb{E}_{\nu_N} [\exp(-\Phi_N)]^{-1} \right\|_{L_{\mu_0}^{p'_1}(\mathcal{U})}, \left\| \exp(\Phi) \right\|_{L_{\mu_0}^{p_1}(\mathcal{U})} \right\} \leq C_1(p_1);$
- (b) $\left\| \mathbb{E}_{\nu_N} \left[\left(\exp(-\Phi) + \exp(-\Phi_N) \right)^{p_2} \right]^{1/p_2} \right\|_{L_{\mu_0}^{2p'_1 p'_3}(\mathcal{U})} \leq C_2(p_1, p_2, p_3);$
- (c) $C_3^{-1} \leq \mathbb{E}_{\nu_N} [Z_N^S] \leq C_3.$

Then there exists $C = C(C_1, C_2, C_3, Z) > 0$, independent of N , such that

$$d_H(\mu, \mu_N^M) \leq C \left\| \mathbb{E}_{\nu_N} [|\Phi - \Phi_N|^{p'_2}]^{1/p'_2} \right\|_{L_{\mu_0}^{2p'_1 p'_3}(\mathcal{U})}, \quad (3.3a)$$

$$C(C_1, C_2, C_3, Z) = \left(\frac{C_1(p_1)}{Z} + C_3 \max\{Z^{-3}, C_3^3\} \right) C_2^2(p_1, p_2, p_3). \quad (3.3b)$$

In the proof of Theorem 3.1, we show that hypothesis (a) arises as an upper bound on the quantity $\|(e^{-\Phi} + \mathbb{E}_{\nu_N}[e^{-\Phi_N}])^{-1}\|_{L_{\mu_0}^{p'_1}(\mathcal{U})}$. In order for the conclusion of Theorem 3.1 to hold, we need the latter to be finite. Thus, hypothesis (a) is an exponential decay condition on the *positive* tails of either Φ or Φ_N , with respect to the appropriate measures. Alternatively, by applying Jensen's inequality to $\mathbb{E}_{\nu_N}[e^{-\Phi_N}]^{-1}$, one can strengthen hypothesis (a) into the hypothesis of exponential integrability of either Φ with respect to μ_0 or Φ_N with respect to $\nu_N \otimes \mu_0$; this yields the same interpretation. Thus, the parameter p_1 quantifies the exponential decay of the positive tail of either Φ or Φ_N .

By comparing the quantity $\|(e^{-\Phi} + \mathbb{E}_{\nu_N}[e^{-\Phi_N}])^{-1}\|_{L_{\mu_0}^{p'_1}(\mathcal{U})}$ from hypothesis (a) with the quantity in hypothesis (b), it follows that hypothesis (b) is an exponential decay condition on the *negative* tails of both Φ and Φ_N . The two new parameters in this decay condition arise because we apply Hölder's inequality twice in order to develop the desired bound (3.3) on $d_H(\mu, \mu_N^M)$. The key desideratum here is that the bound is *multiplicative* in some $L_{\mu_0}^{p'_1}(\mathcal{U})$ -norm of $\mathbb{E}_{\nu_N}[|\Phi - \Phi_N|^{p'_2}]^{1/p'_2}$. The two new parameters p_2 and $p'_1 p'_3$ quantify the decay with respect to ν_N and μ_0 respectively. Note that the interaction between the hypotheses (a) and (b) as described by the conjugate exponent pair (p_1, p'_1) implies that one can trade off faster exponential decay of one tail with slower exponential decay of the other.

The two-sided condition on $\mathbb{E}_{\nu_N}[Z_N^S]$ in hypothesis (c) ensures that both tails of Φ_N with respect to $\nu_N \otimes \mu_0$ decay sufficiently quickly. This hypothesis ensures that the Radon–Nikodym derivative in (3.2) is well-defined.

Finally, we note that the quantity on the right hand side of (3.3a) depends directly on the conjugate exponents of p_1, p_2 and p_3 appearing in hypotheses (a) and (b). The more well behaved the quantities in these hypotheses are, the weaker the norm we can choose on the right hand side of (3.3a).

Theorem 3.2 (Mean-square convergence of the sample posterior). *Suppose that there exist scalars $D_1, D_2 \geq 0$, independent of N , such that, for Hölder-conjugate exponent pairs (q_1, q'_1) and (q_2, q'_2) , we have*

- (a) $\left\| \mathbb{E}_{\nu_N} \left[\left(e^{-\Phi/2} + e^{-\Phi_N/2} \right)^{2q_1} \right]^{1/q_1} \right\|_{L_{\mu_0}^{q'_2}(\mathcal{U})} \leq D_1(q_1, q_2);$
- (b) $\left\| \mathbb{E}_{\nu_N} \left[\left(Z_N^S \max\{Z^{-3}, (Z_N^S)^{-3}\} (e^{-\Phi} + e^{-\Phi_N})^2 \right)^{q_1} \right]^{1/q_1} \right\|_{L_{\mu_0}^{q'_2}(\mathcal{U})} \leq D_2(q_1, q_2).$

Then

$$\mathbb{E}_{\nu_N} \left[d_H(\mu, \mu_N^S)^2 \right]^{1/2} \leq (D_1 + D_2) \left\| \mathbb{E}_{\nu_N} \left[|\Phi - \Phi_N|^{2q_1'} \right]^{1/2q_1'} \right\|_{L_{\mu_0}^{2q_2'}(\mathcal{U})}, \quad (3.4)$$

Hypothesis (a) of Theorem 3.2 arises during the proof as a result of developing an upper bound on $\|\mathbb{E}_{\nu_N}[(e^{-\Phi/2} - e^{-\Phi_N/2})^2]\|$ that is multiplicative in some $L_{\mu_0}^{p'}(\mathcal{U})$ -norm of $\|\mathbb{E}_{\nu_N}[|\Phi - \Phi_N|^{2q_1'}]^{1/2q_1'}\|$. Thus, it describes an exponential decay condition of the negative tails of both Φ or Φ_N ; in particular, hypothesis (a) is always satisfied when the potentials Φ or Φ_N are non-negative, as is usually the case for finite-dimensional data. The appearance of q_1 and q_2 arises due to one application of Hölder's inequality for fulfilling the desideratum of multiplicativity, and q_1 and q_2 quantify the decay with respect to ν_N and μ_0 respectively.

Hypothesis (b) of Theorem 3.2 arises as a result of developing an upper bound on the quantity $\mathbb{E}_{\nu_N}[Z_N^S(Z^{-1/2} - (Z_N^S)^{-1/2})^2]$ that fulfills the desideratum of multiplicativity mentioned above. The presence of both Z_N^S and its reciprocal indicates that hypothesis (b) is analogous to hypothesis (c) of Theorem 3.1, in that hypothesis (b) is a condition on the tails of Φ_N with respect to μ_0 . The difference between hypothesis (b) of Theorem 3.2 and hypothesis (c) of Theorem 3.1 arises due to the fact that the Radon–Nikodym derivative in (3.1) features Z_N^S instead of $\mathbb{E}_{\nu_N}[Z_N^S]$.

We now show that the assumptions of Theorems 3.1 and 3.2 are satisfied when the exact potential Φ and the approximation quality $\Phi_N \approx \Phi$ are suitably well behaved. Since $0 < Z < \infty$, it follows that $C_3^{-1} < Z < C_3$ for some $0 < C_3 < \infty$.

Assumption 3.3. There exists $C_0 \in \mathbb{R}$ that does not depend on N , such that, for all $N \in \mathbb{N}$,

$$\Phi \geq -C_0 \quad \text{and} \quad \nu_N(\{\Phi_N \mid \Phi_N \geq -C_0\}) = 1, \quad (3.5)$$

and for any $0 < C_3 < \infty$ with the property that $C_3^{-1} < Z < C_3$, there exists $N^*(C_3) \in \mathbb{N}$ such that, for all $N \geq N^*$,

$$\mathbb{E}_{\mu_0}[\mathbb{E}_{\nu_N}[|\Phi_N - \Phi|]] \leq \frac{1}{2 \exp(C_0)} \min \left\{ Z - \frac{1}{C_3}, C_3 - Z \right\}. \quad (3.6)$$

The lower bound conditions in (3.5) ensure that the hypothesised exponential decay conditions on the negative tails of the true likelihood and the random likelihoods from Theorems 3.1 and 3.2 are satisfied. The uniform lower bound on Φ translates into a uniform upper bound of the Radon–Nikodym derivative of the posterior with respect to the prior, and is a very mild condition that is satisfied in many, if not most, BIPs. Given this fact, it is reasonable to demand that the Φ_N satisfy the same uniform lower bound, ν_N -almost surely and for all $N \in \mathbb{N}$; this is the content of the second condition in (3.5). Condition (3.6) expresses the condition that, by choosing N sufficiently large, one can approximate Φ arbitrarily well using the random Φ_N , with respect to the $L_{\mu_0 \otimes \nu_N}^1$ topology. This assumption ensures that the stated aims of this work are reasonable.

Lemma 3.4. Suppose that Assumption 3.3 holds with C_0 as in (3.5) and C_3 and $N^*(C_3)$ as in (3.6), that $\exp(\Phi) \in L_{\mu_0}^{p^*}(\mathcal{U})$ for some $1 \leq p^* \leq +\infty$ with conjugate exponent $(p^*)'$, and there exists some $C_4 \in \mathbb{R}$ that does not depend on N , such that, for all $N \in \mathbb{N}$,

$$\nu_N(\{\Phi_N \mid \mathbb{E}_{\mu_0}[\Phi_N] \leq C_4\}) = 1. \quad (3.7)$$

Then the hypotheses of Theorem 3.1 hold, with

$$p_1 = p^*, \quad p_2 = p_3 = +\infty, \quad C_1 = \|\exp(\Phi)\|_{L_{\mu_0}^{p^*}(\mathcal{U})}, \quad C_2 = 2 \exp(C_0),$$

and C_3 as above. Moreover, the hypotheses of Theorem 3.2 hold, with

$$q_1 = q_2 = \infty, \quad D_1 = 4 \exp(C_0), \quad D_2 = 4 \exp(3C_0) \max\{C_3^{-3}, \exp(3C_4)\}.$$

The uniform upper bound condition on Φ_N with respect to μ_0 in (3.7) is rather strong; we use it to ensure that Z_N^S is bounded away from zero, uniformly with respect to Φ_N and $N \in \mathbb{N}$. Together with the condition on Φ_N in (3.5), this translates to uniform lower and upper bounds on Z_N^S ; the latter implies that hypothesis (b) in Theorem 3.2 holds with the stated values of q_1 and q_2 . A sufficient condition for (3.7) is that the Φ_N are themselves uniformly bounded. This condition is of interest when the misfit Φ is associated to a bounded forward model and the data take values in a bounded subset.

Lemma 3.5. *Suppose that Assumption 3.3 holds with C_0 as in (3.5) and C_3 and $N^*(C_3)$ as in (3.6), and that there exists some $2 < \rho^* < +\infty$ such that $\mathbb{E}_{\nu_N}[\exp(\rho^* \Phi_N)] \in L_{\mu_0}^1(\mathcal{U})$. Then the hypotheses of Theorem 3.1 hold, with*

$$p_1 = \rho^*, \quad p_2 = p_3 = +\infty, \quad C_1 = \|\mathbb{E}_{\nu_N}[\exp(\rho^* \Phi_N)]\|_{L_{\mu_0}^1(\mathcal{U})}^{1/\rho^*}, \quad C_2 = 2 \exp(C_0),$$

and C_3 as above. Moreover, the hypotheses of Theorem 3.2 hold, with

$$\begin{aligned} q_1 &= \frac{\rho^*}{2}, & q_2 &= +\infty, \\ D_1 &= 4 \exp(C_0), & D_2 &= 4 \exp(2C_0) \left(C_3^{-3} \exp(C_0) + \|\mathbb{E}_{\nu_N}[\exp(\rho^* \Phi_N)]\|_{L_{\mu_0}^1(\mathcal{U})}^{2/\rho^*} \right). \end{aligned}$$

By comparing the hypotheses and conclusions of Lemma 3.4 and Lemma 3.5, we observe that, by reducing the exponent of integrability from $q_1 = +\infty$ to $q_1 = \rho^*/2$, we can replace the strong uniform upper bound condition (3.7) on Φ_N from Lemma 3.4 with the weaker condition that $\exp(\Phi_N) \in L_{\mu_0}^{\rho^*}(\mathcal{U})$ in Lemma 3.5, and thus increase the scope of applicability of the conclusion.

In Lemmas 3.4 and 3.5 above, we have specified the largest possible values of the exponents that are compatible with the hypotheses. This is because later, in Theorem 3.9, we will want to use the smallest possible values of the corresponding *conjugate* exponents in the resulting inequalities (3.3a) and (3.4).

3.2 Random forward models in quadratic potentials

In many settings, the potentials Φ and Φ_N have a common form and differ only in the parameter-to-observable map. In this section we shall assume that Φ and Φ_N are quadratic misfits of the form

$$\Phi(u) = \frac{1}{2} \|\Gamma^{-1/2}(G(u) - y)\|^2 \quad \text{and} \quad \Phi_N(u) = \frac{1}{2} \|\Gamma^{-1/2}(G_N(u) - y)\|^2, \quad (3.8)$$

corresponding to centred Gaussian observational noise with symmetric positive-definite covariance Γ . Again, we assume that G is deterministic while G_N is random. In this section, for this setting, we show how the quality of the approximation $G_N \approx G$ transfers to the approximation $\Phi_N \approx \Phi$, and hence to the approximation $\mu_N \approx \mu$ (for either the sample or marginal approximate posterior).

Pointwise in u and ω , the errors in the misfit and the forward model are related according to the following proposition.

Proposition 3.6. *Let Φ and Φ_N be defined as in (3.8), where $\mathcal{Y} = \mathbb{R}^J$ for some $J \in \mathbb{N}$ and the eigenvalues of the operator Γ are bounded away from zero. Then, for some $C = C_\Gamma > 0$, for all $u \in \mathcal{U}$, and ν_N -almost surely*

$$|\Phi(u) - \Phi_N(u)| \leq 2C_\Gamma \left(\Phi(u)^{1/2} \|G(u) - G_N(u)\| + \|G(u) - G_N(u)\|^2 \right). \quad (3.9)$$

Hence, for $q \in [1, \infty)$ and all $u \in \mathcal{U}$,

$$\begin{aligned} \mathbb{E}_{\nu_N} [|\Phi(u) - \Phi_N(u)|^q]^{1/q} &\leq 4C_\Gamma \left(\Phi(u)^{q/2} \mathbb{E}_{\nu_N} [\|G(u) - G_N(u)\|^q] \right. \\ &\quad \left. + \mathbb{E}_{\nu_N} [\|G(u) - G_N(u)\|^{2q}] \right)^{1/q}. \end{aligned} \quad (3.10)$$

By assuming that $\mathcal{Y} = \mathbb{R}^J$, we assume that the data live in a finite-dimensional space. This is a standard assumption in the area, and implies that the operator Γ is simply a matrix. The assumption of the eigenvalues of Γ being bounded away from zero is equivalent to assuming that Γ is invertible, which follows immediately from the assumption stated earlier that Γ is a symmetric and positive-definite covariance matrix.

Corollary 3.7. *Let $1 \leq q \leq s$, and suppose that $\Phi \in L_{\mu_0}^s(\mathcal{U})$. If there exists an $N^* \in \mathbb{N}$ such that, for all $N \geq N^*$,*

$$\left\| \mathbb{E}_{\nu_N} [\|G - G_N\|^{2q}]^{1/q} \right\|_{L_{\mu_0}^s(\mathcal{U})} \leq 1,$$

then, there exists some $C = C(s) > 0$ that does not depend on N such that for all $N \geq N^*$,

$$\left\| \mathbb{E}_{\nu_N} [|\Phi - \Phi_N|^q]^{1/q} \right\|_{L_{\mu_0}^s(\mathcal{U})} \leq C \left\| \mathbb{E}_{\nu_N} [\|G - G_N\|^{2q}]^{1/q} \right\|_{L_{\mu_0}^s(\mathcal{U})}^{1/2}$$

where $C(s) = (8C_\Gamma)(\mathbb{E}_{\mu_0}[\Phi^s]^{1/2} + 1)^{1/s}$ and C_Γ is as in Proposition 3.6.

The hypotheses ensure that the integrability of the misfit Φ determines the highest degree of integrability of the forward operators G_N and G , and that for sufficiently large N , we may make the norm of the difference of $G - G_N$ in an appropriate topology small enough. The constraint (3.7) is used to combine the $\|G(u) - G_N(u)\|$ and $\|G(u) - G_N(u)\|^2$ terms in (3.9). The resulting simplification ensures that we may apply Lemma 3.8.

Lemma 3.8. *Let Φ and Φ_N be as in (3.8). If, for some $q, s \geq 1$,*

$$\lim_{N \rightarrow \infty} \left\| \mathbb{E}_{\nu_N} [\|G - G_N\|^{2q}]^{1/q} \right\|_{L_{\mu_0}^s(\mathcal{U})} = 0, \quad (3.11)$$

then Assumption 3.3 holds.

The lemma states that if the random forward model converges to the true forward model in the appropriate topology, then the conditions in Assumption 3.3 are satisfied by the corresponding random misfits. Since the misfits were assumed to be quadratic in (3.8), the key contribution of Lemma 3.8 is to ensure that the approximation quality condition (3.6) is satisfied.

We shall use the preceding results to obtain bounds on the Hellinger distance in terms of errors in the forward model, of the following form: for $C, D > 0$ and $r_1, r_2, s_1, s_2 \geq 1$ that do not depend on N ,

$$d_H(\mu, \mu_N^M) \leq C \left\| \mathbb{E}_{\nu_N} [\|G_N - G\|^{2r_1}]^{1/r_1} \right\|_{L_{\mu_0}^{r_2}(\mathcal{U})}^{1/2} \quad (3.12)$$

$$\mathbb{E}_{\nu_N} \left[d_H(\mu, \mu_N^S)^2 \right]^{1/2} \leq D \left\| \mathbb{E}_{\nu_N} [\|G_N - G\|^{2s_1}]^{1/s_1} \right\|_{L_{\mu_0}^{s_2}(\mathcal{U})}^{1/2}. \quad (3.13)$$

For brevity and simplicity, the following result uses one pair $q, s \geq 1$ in (3.11) in order to obtain convergence statements for *both* μ_N^M and μ_N^S . If one is interested in only one of these measures, then one may optimise q and s accordingly.

Theorem 3.9 (Convergence of posteriors for randomised forward models in quadratic potentials). *Let Φ and Φ_N be as in (3.8).*

(a) *Suppose there exists some $p^* > 1$ with Hölder conjugate $(p^*)'$ such that $\exp(\Phi) \in L_{\mu_0}^{p^*}(\mathcal{U})$, and suppose that (3.7) holds for some $C_4 > 0$. If $G_N \rightarrow G$ as in (3.11) with $q = 2$ and $s = 2p^*/(p^* - 1)$, then the following hold:*

- (i) *there exists some $C > 0$ that does not depend on N , for which (3.12) holds with $r_1 = 1$ and $r_2 = 2p^*/(p^* - 1)$, and*
- (ii) *there exists some $D > 0$ that does not depend on N , for which (3.13) holds with $s_1 = 2$ and $s_2 = 2$.*

(b) *Suppose there exists some $2 < \rho^* < \infty$ such that $\mathbb{E}_{\nu_N}[\exp(\rho^* \Phi_N)] \in L_{\mu_0}^1$. If $G_N \rightarrow G$ as in (3.11) with $q = 2\rho^*/(\rho^* - 2)$ and $s = 2\rho^*/(\rho^* - 1)$, then the following hold:*

- (i) *there exists some $C > 0$ that does not depend on N , for which (3.12) holds with $r_1 = 1$ and $r_2 = 2\rho^*/(\rho^* - 1)$, and*
- (ii) *there exists some $D > 0$ that does not depend on N , for which (3.13) holds with $s_1 = 2\rho^*/(\rho^* - 2)$ and $s_2 = 2$.*

In both cases, μ_N^M and μ_N^S converge to μ in the appropriate metrics given in (3.12) and (3.13) respectively.

The proof of Theorem 3.9 consists of tracking the dependence of the parameters over the sequential application of the preceding results, all of which are used.

Case (a) applies in the situation where the random approximations Φ_N are uniformly bounded from above; as discussed earlier, this condition is satisfied in the case that the misfit Φ is associated to a bounded forward model and the data take values in a bounded subset of $\mathcal{Y} = \mathbb{R}^J$. Note that the

topology of the convergence of G_N to G is quantified by s and q , and that s depends on the parameter p^* that quantifies the exponential μ_0 -integrability of the misfit Φ . In particular, the faster the exponential decay of the positive tail of Φ (i.e. the larger the value of p^*), the stronger the topology of convergence of G_N to G .

In contrast to case (a), case (b) does not assume that the misfit Φ is exponentially integrable or that the random approximations Φ_N are uniformly bounded from above ν_N -almost surely. Instead, exponential integrability of the random misfit Φ_N is required. Another difference is that the exponential integrability parameter ρ^* determines the strength of the topology of convergence of the random forward models, not only with respect to the μ_0 -topology, but also to the ν_N -topology as well.

4 Application: randomised misfit models

This section considers a particular Monte Carlo approximation Φ_N of a quadratic potential Φ , proposed by Nemirovski et al. (2008); Shapiro et al. (2009), and further applied and analysed in the context of BIPs by Le et al. (2017). This approximation is particularly useful when the data $y \in \mathbb{R}^J$ has very high dimension, so that one does not wish to interrogate every component of the data vector y , or evaluate every component of the model prediction $G(u)$ and compare it with the corresponding component of y .

Let σ be an $\mathbb{R}^J$ -valued random vector with mean zero and identity covariance, and let $\sigma^{(1)}, \dots, \sigma^{(N)}$ be independent and identically distributed copies (samples) of σ . We then have the following approximation:

$$\begin{aligned}\Phi(u) &:= \frac{1}{2} \left\| \Gamma^{-1/2}(y - G(u)) \right\|^2 \\ &= \frac{1}{2} (\Gamma^{-1/2}(y - G(u)))^T \mathbb{E}[\sigma \sigma^T] (\Gamma^{-1/2}(y - G(u))) \\ &= \frac{1}{2} \mathbb{E} \left[\left| \sigma^T (\Gamma^{-1/2}(y - G(u))) \right|^2 \right] \\ &\approx \frac{1}{2N} \sum_{i=1}^N \left| \sigma^{(i)T} (\Gamma^{-1/2}(y - G(u))) \right|^2 \\ &=: \Phi_N(u).\end{aligned}$$

The analysis and numerical studies in Le et al. (2017, Sections 3–4) suggest that a good choice for the random vector σ would be one with independent and identically distributed (i.i.d.) entries from a sub-Gaussian probability distribution on $\mathbb{R}$. Examples of sub-Gaussian distributions considered include

- (a) the standard Gaussian distribution: $\sigma_j \sim \mathcal{N}(0, 1)$, for $j = 1, \dots, J$; and
- (b) the ℓ -sparse distribution: for $\ell \in [0, 1)$, let $s := \frac{1}{1-\ell} \geq 1$ and set, for $j = 1, \dots, J$,

$$\sigma_j := \sqrt{s} \begin{cases} 1, & \text{with probability } \frac{1}{2s}, \\ 0, & \text{with probability } \ell = 1 - \frac{1}{s}, \\ -1, & \text{with probability } \frac{1}{2s}. \end{cases}$$

The randomised misfit Φ_N can provide computational benefits in two ways. Firstly, a single evaluation of Φ_N can be made cheap by choosing the ℓ -sparse distribution for σ , with large sparsity parameter ℓ . This choice ensures that a large proportion of the entries of each sample $\sigma^{(i)}$ will be zero, significantly reducing the cost to compute the required inner products in Φ_N , since there is no need to compute the components of the data or model vector that will be eliminated by the sparsity pattern. The value of N of course also influences the computational cost. It is observed by Le et al. (2017) that, for large J and moderate $N \approx 10$, the random potential Φ_N and the original potential Φ are already very similar, in particular having approximately the same minimisers and minimum values. Statistically, these correspond to the maximum likelihood estimators under Φ and Φ_N being very similar; after weighting by a prior, this corresponds to similarity of maximum a posteriori (MAP) estimators.

The second benefit of the randomised misfit approach, and the main motivation for its use in Le et al. (2017), is the reduction in computational effort needed to compute the MAP estimate. This task involves the solution of a large-scale optimisation problem involving Φ in the objective function, which is

typically done using inexact Newton methods. It is shown by Le et al. (2017) that the required number of evaluations of the forward model G and its adjoint is drastically reduced when using the randomised misfit Φ_N as opposed to using the true misfit Φ , approximately by a factor of $\frac{J}{N}$.

The aim of this section is to show that the use of the randomised misfit Φ_N does not only lead to the MAP estimate being well-approximated, but in fact the whole Bayesian posterior distribution. Thus, the corresponding conjecture is that the ideal and deterministic posterior $d\mu(u) \propto \exp(-\Phi(u)) d\mu_0(u)$ is well approximated by the random posterior $d\mu_N^S(u) \propto \exp(-\Phi_N(u)) d\mu_0(u)$. Indeed, via Theorem 3.2, we have the following convergence result for the case of a sparsifying distribution:

Proposition 4.1. *Suppose that the entries of σ are i.i.d. ℓ -sparse, for some $\ell \in [0, 1)$, and that $\Phi \in L_{\mu_0}^2(\mathcal{U})$. Then there exists a constant C , independent of N , such that*

$$\left(\mathbb{E}_{\nu_N} [d_H(\mu, \mu_N^S)^2] \right)^{1/2} \leq \frac{C}{\sqrt{N}}. \quad (4.1)$$

(In this section, ν_N plays the role of the distribution of $\sigma^{(1)}, \dots, \sigma^{(N)}$.) As the proof reveals, a valid choice of the constant C in (4.1) is

$$C = (D_1 + D_2) \sqrt{J^3 \mathbb{E}_{\nu_N} [\sigma_j^4] - 1} \|\Phi\|_{L_{\mu_0}^2(\mathcal{U})} = (D_1 + D_2) \sqrt{J^3 s^3 - 1} \|\Phi\|_{L_{\mu_0}^2(\mathcal{U})}, \quad (4.2)$$

where the constant $(D_1 + D_2)$ is as in Theorem 3.2. Thus, as one would expect, the accuracy of the approximation decreases as σ approaches the complete sparsification case $\ell = 1$ or as the data dimension J increases, but always with the same convergence rate $N^{-1/2}$ in terms of the approximation dimension N .

Remark 4.2. The proof of Proposition 4.1 can be modified to yield the same result for arbitrary i.i.d. σ_j with bounded support, though the sparsifying case is obviously the one with the easiest interpretation. However, extending Proposition 4.1 to the case of i.i.d. Gaussian random variables $\sigma_j \sim \mathcal{N}(0, 1)$ appears to be problematic. In the proof, we crucially make use of the bound $|\sigma_j| \leq \sqrt{s}$ to verify Assumption (b) of Theorem 3.2. For Gaussian random variables, we would similarly need an N -independent bound on the exponential moments of

$$\max_{\substack{1 \leq i \leq N \\ 1 \leq j \leq J}} \sigma_j^{(i)},$$

which is not possible. We leave this as an interesting question for future work: would a different proof strategy yield convergence in the Gaussian case, or is the Gaussian setting genuinely one in which the MAP problem is well approximated but the BIP is not?

5 Application: probabilistic integration of dynamical systems

The data-based inference of initial conditions or governing parameters for dynamical problems arises frequently in scientific applications, a prime example being data assimilation in numerical weather prediction (Law et al., 2015; Reich and Cotter, 2015). In this setting, the Bayesian likelihood involves a solution of the mathematical model for the dynamics, which is typically an ODE or time-dependent PDE; we focus here on the ODE situation. Even when the governing ODE is deterministic, it may be profitable to perform a probabilistic numerical solution: possible motivations for doing so include the representation of model error (model inadequacy) in the ODE itself, and the impact of discretisation uncertainty. When such a probabilistic solver is used for the ODE, the likelihood becomes random in the sense considered in this paper.

Random approximate solution of deterministic ODEs is an old idea (Diaconis, 1988; Skilling, 1992) that has received renewed attention in recent years (Conrad et al., 2016; Hennig et al., 2015; Lie et al., 2017; Schober et al., 2014). As random forward models, these probabilistic ODE solvers are amenable to the analysis of Section 3. Let $f: \mathbb{R}^d \rightarrow \mathbb{R}^d$ and consider the following parameter-dependent initial value problem for a fixed, parameter-independent duration $T > 0$:

$$\begin{aligned} \frac{d}{dt} z(t; u) &= f(z(t; u); u), & \text{for } 0 \leq t \leq T, \\ z(0; u) &= z_0(u). \end{aligned} \quad (5.1)$$

In the context of the BIP presented in Section 2, the unknown parameter u will appear in the definition of the initial condition $z_0 = z_0(u)$ or the right-hand side $f(z(t)) = f(z(t); u)$, resulting in the parameter-dependent solution $(z(t; u))_{t \in [0, T]}$. Define the solution operator

$$S: \mathcal{U} \rightarrow C([0, T]; \mathbb{R}^d), \quad u \mapsto S(u) := (z(t; u))_{t \in [0, T]}, \quad (5.2)$$

where $(z(t; u))_{t \in [0, T]}$ solves (5.1). We equip $C([0, T]; \mathbb{R}^d)$ with the supremum norm.

For notational convenience, we will for the majority of this section not indicate the dependence of z_0 or f on u . We will, however, explicitly track the dependence on z_0 and f of the error analysis below.

Let $F^t: \mathbb{R}^d \rightarrow \mathbb{R}^d$ be the flow map associated to the initial value problem (5.1), i.e. $F^t(z_0) := z(t; u) = S(u)(t)$. Fix a time step $\tau > 0$ such that $N := T/\tau \in \mathbb{N}$, and a time grid

$$t_k := k\tau \text{ for } k \in [N] := \{0, 1, \dots, N\}. \quad (5.3)$$

We denote by $z_k := z(t_k) \equiv F^\tau(z_{k-1})$ the value of the exact solution to (5.1) at time t_k . We shall sometimes abuse notation and write $[N] = \{0, 1, \dots, N-1\}$ or $[N] = \{1, 2, \dots, N\}$.

To a single-step numerical integration method (e.g. a Runge–Kutta method of some order) we shall associate a numerical flow map $\Psi^\tau: \mathbb{R}^d \rightarrow \mathbb{R}^d$. The numerical flow map approximates the sequence $(z_k)_{k \in [N]}$ by a sequence $(Z'_k)_{k \in [N]}$, where $Z'_k := \Psi^\tau(Z'_{k-1})$. A fundamental task in numerical analysis is to determine sufficient conditions for convergence of the sequence $(Z'_k)_{k \in [N]}$ to $(z_k)_{k \in [N]}$. The investigations of Conrad et al. (2016) and Lie et al. (2017) concern a similar task in the context of uncertainty quantification. Given $\tau > 0$, consider a collection $(\xi_k)_{k \in [N]}$ of stochastic processes $\xi_k: \Omega \times [0, \tau] \rightarrow \mathbb{R}^d$ having almost-surely continuous paths. Define a stochastic process $(Z_t)_{t \in [0, T]}$ in terms of a new randomised integrator

$$Z(t_{k+1}; u) := \Psi^\tau(Z(t_k; u)) + \xi_k(\tau). \quad (5.4)$$

The stochastic processes $(\xi_k)_{k \in [N]}$ are intended to capture the effect of uncertainties, e.g. those that arise due to properties of the vector field that are not resolved by the time grid (5.3) associated to the time step τ . We extend the definition (5.4) to continuous time via

$$Z(t; u) := \Psi^{t-t_k}(Z(t_k; u)) + \xi_k(t - t_k), \quad \text{for } t_k < t < t_{k+1}. \quad (5.5)$$

We shall use the $(\xi_k)_{k \in [N]}$ to construct our random approximations to Φ . Note therefore that, in order to be consistent with our assumption (see the third paragraph of Section 3) that the randomness in the approximation of Φ is independent of the randomness in the parameter u being inferred, we shall assume that the $(\xi_k)_{k \in [N]}$ do not depend on the parameter u . However, the map Ψ^τ does depend on the parameter $u \in \mathcal{U}$, because Ψ^τ involves the vector field $f(\cdot; u)$.

Define the random solution operator associated to the randomised integrator (5.5):

$$S_N: \mathcal{U} \rightarrow C([0, T]; \mathbb{R}^d), \quad u \mapsto S_N(u) := (Z(t; u))_{t \in [0, T]}, \quad (5.6)$$

where $(Z(t; u))_{t \in [0, T]}$ satisfies (5.5), and is almost surely continuous.

Let $T_J \subset [0, T]$ be a strictly increasing sequence of time points, indexed by a finite, nonempty index set J with cardinality $|J| \in \mathbb{N}$. Note that T_J may coincide with the time grid defined in (5.3); to increase the scope of the subsequent analysis however, we allow for T_J to differ from (5.3). Let $\mathcal{Y} := \mathbb{R}^{|J|}$, and equip it with the topology induced by the standard Euclidean inner product. Define the observation operator

$$O: C([0, T]; \mathbb{R}^d) \rightarrow \mathcal{Y}, \quad \tilde{z} \mapsto O(\tilde{z}) := (\tilde{z}(t_j))_{t_j \in T_J}, \quad (5.7)$$

which projects some $\tilde{z} \in C([0, T]; \mathbb{R}^d)$ to a finite-dimensional vector in $\mathcal{Y}$ constructed by stacking the $\mathbb{R}^d$ -valued vectors that result from evaluating $\tilde{z}$ at the time points in T_J . We take the norm on $\mathcal{Y}$ to be $\|\cdot\|_{\ell_2^{|J|}}$.

Given the operators S , O , and S_N defined in (5.2), (5.7), and (5.6), we define the forward operators $G, G_N: \mathcal{U} \rightarrow \mathcal{Y}$ by

$$G := O \circ S, \quad G_N := O \circ S_N. \quad (5.8)$$

The associated likelihoods are the quadratic misfits given by (3.8) with some fixed, positive-definite matrix Γ .

We define the continuous-time error process by

$$e(t; u) := z(t; u) - Z(t; u), \quad 0 \leq t \leq T. \quad (5.9)$$

Since T_J is a proper subset of $[0, T]$, it follows that

$$\|G_N(u) - G(u)\| \leq |J| \sup_{0 \leq t \leq T} \|e(t; u)\|_{\ell_2^d}. \quad (5.10)$$

This completes our formulation of the probabilistic numerical integration of the ODE (5.1) as a random likelihood model of the type considered in Section 3.

5.1 Convergence in continuous time for Lipschitz flows

In this section, we quote some assumptions and results from Lie et al. (2017). The vector field f in (5.1) induces a flow $F^\tau: \mathbb{R}^d \rightarrow \mathbb{R}^d$ by

$$F^\tau(a) = a + \int_0^\tau f(F^t(a)) dt. \quad (5.11)$$

Assumption 5.1 (Assumption 3.1, Lie et al. (2017)). The vector field f admits $0 < \tau^* \leq 1$ and $C_F \geq 1$, such that for $0 < \tau < \tau^*$, the flow $F^\tau: \mathbb{R}^d \rightarrow \mathbb{R}^d$ defined by (5.11) is globally Lipschitz, with

$$\|F^\tau(z_0) - F^\tau(v_0)\| \leq (1 + C_F \tau) \|z_0 - v_0\|, \quad \text{for all } z_0, v_0 \in \mathbb{R}^d.$$

A globally Lipschitz vector field f in (5.1) yields a flow map F^t that satisfies Assumption 5.1. However, vector fields that satisfy a one-sided Lipschitz condition also have the same property. Such vector fields have been studied in the numerical analysis literature for both ordinary and stochastic differential equations in the last four decades; see, e.g. (Butcher, 1975), and the references cited in Section 3.1 of (Higham et al., 2002).

Recall that $\Psi^\tau: \mathbb{R}^d \rightarrow \mathbb{R}^d$ represents the numerical method that we use to integrate (5.1).

Assumption 5.2 (Assumption 3.2, Lie et al. (2017)). The numerical method Ψ^τ has uniform local truncation error of order $q + 1$: for some constant $C_\Psi \geq 1$ that does not depend on τ ,

$$\sup_{v \in \mathbb{R}^d} \|\Psi^\tau(v) - F^\tau(v)\| \leq C_\Psi \tau^{q+1}.$$

The assumption above is satisfied for both single-step and multistep numerical methods that are obtained by considering vector fields in $C^q(\mathbb{R}^d)$, provided that the q^{th} derivatives are bounded; see Section III.2 of (Hairer et al., 2009). We emphasise that the above assumption is made to simplify the analysis, and that the convergence results below extend to the case where the uniform bound does not hold; see Section 4 of (Lie et al., 2017).

Now recall the collection $(\xi_k(\tau))_{k \in [N]}$ of random variables, where $\xi_k(\tau)$ is used in (5.4).

Assumption 5.3 (Assumption 5.1, Lie et al. (2017)). The stochastic processes $(\xi_k)_{k \in \mathbb{N}}$ admit $p \geq 1$, $R \in \mathbb{N} \cup \{+\infty\}$, and $C_{\xi, R} \geq 1$, independent of k and τ , such that for all $1 \leq r \leq R$ and all $k \in \mathbb{N}$,

$$\mathbb{E}_{\nu_N} \left[\sup_{0 \leq t \leq T/N} \|\xi_k(t)\|^r \right] \leq \left(C_{\xi, R} \left(\frac{T}{N} \right)^{p+1/2} \right)^r.$$

(In this section, ν_N plays the role of the distribution of the ξ_k 's.) The assumption above quantifies the regularity of the $(\xi_k)_{k \in [N]}$ by specifying how many moments each $\xi_k(t)$ has and how quickly these decay with $\tau = T/N$. We do not require the $(\xi_k(t))_{k \in [N]}$ to have zero mean, to be independent, or to be identically distributed. It is shown in Section 5.2 of (Lie et al., 2017) that, for example, the integrated Brownian motion process satisfies Assumption 5.3. The integrated Brownian motion process has been used as a state-independent model of the uncertainty in the off-grid behaviour of solutions to ODEs in (Conrad et al., 2016; Schober et al., 2014; Chkrebtii et al., 2016).

We now consider the following convergence theorem:

Theorem 5.4 (Theorem 5.2, Lie et al. (2017)). *Suppose that $e_0 = 0$, and suppose that Assumptions 5.1, 5.2, and 5.3 hold with parameters τ^* , C_F , C_Ψ , q , $C_{\xi,R}$, p , and R . Let $n \in \mathbb{N}$, with $n \leq R$. Then, for all $T/\tau^* < N$,*

$$\mathbb{E}_{\nu_N} \left[\sup_{0 \leq t \leq T} \|e(t; u)\|^n \right] \leq 3^{n-1} ((1 + C_F \tau^*)^n \overline{C} + C_\Psi^n (\tau^*)^n + T C_{\xi,R}^n) \left(\frac{T}{N} \right)^{n(q \wedge (p-1/2))}, \quad (5.12)$$

where

$$\begin{aligned} \overline{C} &:= 2T \max\{(4C_\Psi)^n, (2C_{\xi,R})^n\} \exp(T C_F(n, \tau^*)) \\ C_F(n, \tau^*) &:= [(1 + \tau^* 2^{n-1})^2 (1 + \tau^* C_F)^n - 1] (\tau^*)^{-1}. \end{aligned}$$

Note that the scalars $\overline{C}$ and $C_F(n, \tau^*)$ depend on $u \in \mathcal{U}$, since C_F and C_Ψ depend on the vector field f , which in turn depends on the parameter u .

Recall that the random variable $(Z(t; u))_{0 \leq t \leq T}$ defined in (5.5) is a random surrogate for the true solution of the ODE (5.1). The law of $(Z(t; u))_{0 \leq t \leq T}$ is thus a probability measure on the space of continuous paths defined on the interval $[0, T]$. With this in mind, the interpretation of Theorem 5.4 is that the law of $(Z(t; u))_{0 \leq t \leq T}$ contracts to the Dirac distribution located at the true solution $(z(t; u))_{0 \leq t \leq T}$ of (5.1), as the spacing T/N in the time grid (5.3) decreases to zero. Equivalently, given the true solution operator S and its random counterpart S_N , Theorem 5.4 implies that the random solution operator converges in the L^n topology to the true solution operator. Thus, Theorem 5.4 guarantees that by refining the time grid, one reduces the uncertainty over the solution of (5.1). This is a desirable feature for uncertainty quantification, since estimates of the solution uncertainty can also be fed forward to obtain estimates of the uncertainty of *functionals* of the solution, and since the probabilistic description allows for a more nuanced description of the uncertainty compared to the usual worst-case description that is common in the numerical analysis of deterministic methods.

Corollary 5.5 (Corollary 5.3, Lie et al. (2017)). *Fix $n \in \mathbb{N}$. Suppose that Assumptions 5.1 and 5.2 hold, and that Assumption 5.3 holds with $R = +\infty$ and $p \geq 1/2$. Then, for all $0 < \tau < \tau^*$,*

$$\mathbb{E}_{\nu_N} \left[\exp \left(\rho \sup_{0 \leq t \leq T} \|e(t)\|^n \right) \right] < \infty, \quad \text{for all } \rho \in \mathbb{R}. \quad (5.13)$$

Since the exponential integrability of a random variable is related to the exponential concentration of its values about its mean or median, the above result shows that strong assumptions on the model of uncertainty translate to strong conclusions about the behaviour of the corresponding error. In the context of random approximations of BIPs, we shall use Corollary 5.5 in order to establish the convergence of the random approximations in the Hellinger sense.

5.2 Effect of probabilistic integration on Bayesian posterior distribution

Define the approximate posteriors μ_N^M and μ_N^S according to (3.2) and (3.1), using the quadratic misfits Φ and Φ_N from (3.8) and the forward models $G = O \circ S$ and $G_N = O \circ S_N$ given in (5.2) and (5.6) respectively, where O denotes the observation operator associated to a fixed, finite sequence T_J of observation times in $[0, T]$.

As we saw in the last section, the results of Lie et al. (2017) guarantee convergence in the L^n topology of the random solution operator S_N to the true solution operator S . It is of interest to determine whether one can use this result to guarantee that one can perform inference over u using the probabilistic integrator. In particular, given that the probabilistic integrator provides a random approximation, it is of interest to determine whether one can obtain results that are not only reasonable, but that improve as the time resolution of the probabilistic integrator increases. The following result shows that this is indeed the case: as the time resolution increases, the random forward model G_N yields a random posterior over the parameter space that converges in the Hellinger topology to the true posterior at the expected rate.

Theorem 5.6. *Suppose that $\mathcal{U}$ is a compact subset of $\mathbb{R}^m$ for some $m \in \mathbb{N}$, and suppose that $S, S_N: \mathcal{U} \rightarrow C([0, T]; \mathbb{R}^d)$ are continuous maps. Let $2 < \rho^* < \infty$ be arbitrary. Suppose that $e_0 = 0$, and that Assumptions 5.1, 5.2, and 5.3 hold with parameters τ^* , C_F , C_Ψ , q , $R = +\infty$, $C_{\xi,R}$, and p , and that these parameters depend continuously on u . Then, for $N \in \mathbb{N}$ such that $T/\tau^* < N$, the following hold:*

- (a) *there exists some $C > 0$ that does not depend on N , such that (3.12) holds for $r_1 = 1$ and $r_2 = 2\rho^*/(\rho^* - 1)$, and*
- (b) *there exists some $D > 0$ that does not depend on N , such that (3.13) holds for $s_1 = 2\rho^*/(\rho^* - 2)$ and $s_2 = 2$.*

The parameter ρ^* above plays the same role as ρ^* in Theorem 3.9, case (b); ρ^* quantifies the exponential decay of the random misfit Φ_N with respect to ν_N . For this reason, ρ^* is constrained to the same range of values $2 < \rho^* < \infty$ as given there, and determines the parameters r_2 and s_1 that partly describe the convergence rates in (3.12) and (3.13). As shown in the proof, the reason why ρ^* does not appear to play any further role is due to (5.10) and Corollary 5.5. In particular, since Assumption 5.3 holds with $R = +\infty$, Corollary 5.5 ensures that the exponential decay parameter ρ^* need not be constrained to any bounded interval.

The continuous dependence on u of the parameters of Assumptions 5.1, 5.2 and 5.3 also allows for parameters that do not depend on u , e.g. $R = +\infty$. The assumption of continuous dependence on u of the parameters, and the assumptions on $\mathcal{U}$, ensure that the map $u \mapsto \mathbb{E}_{\nu_N}[\exp(\rho^*\Phi_N(u))]$ is uniformly bounded by a scalar that depends only on $\mathcal{U}$; from this the exponential integrability hypothesis on Φ_N of Theorem 3.9(b) holds, and we can apply the corresponding conclusions. While these assumptions may appear to be strong, they simplify the analysis considerably and thus are not uncommon in the literature on parameter inference for dynamical systems. We leave the investigation of weaker assumptions for future work.

6 Concluding remarks

In this paper we have considered the impact upon a BIP of replacing the log-likelihood function Φ by a random function Φ_N . Such approximations occur for example when a cheap stochastic emulator is used in place of an expensive exact log-likelihood, or when a probabilistic solver is used to simulate the forward model.

Our results show that such approximations are well-posed, with the approximate Bayesian posterior distribution converging to the true Bayesian posterior as the error between Φ and Φ_N , measured in a suitable sense, goes to zero. More precisely, we have shown that the convergence rate of the random log-likelihood Φ_N to Φ — as assessed in a nested L^p norm with respect to the distribution ν_N of Φ_N and the Bayesian prior distribution μ_0 of the unknown u — transfers to convergence of two natural approximations to the exact Bayesian posterior μ , namely (a) the randomised posterior measure μ_N^S that simply has Φ_N in place of Φ , and (b) the deterministic pseudo-marginal posterior measure μ_N^M , in which the likelihood function and marginal likelihood of μ_N^S are individually averaged with respect to ν_N .

Since the hypotheses that are required for these results operate directly at the level of finite-order moments of the error $\Phi_N - \Phi$, the convergence results in this paper automatically apply to GP approximations, as previously considered by Stuart and Teckentrup (2017), which have moments of all orders. However, in a substantial generalisation, Theorems 3.1 and 3.2 show that, in the L^2 case,

$$d_H(\mu, \mu_N^M) \leq C \left\| \mathbb{E}_{\nu_N} [|\Phi - \Phi_N|] \right\|_{L_{\mu_0}^2(\mathcal{U})},$$

$$\mathbb{E}_{\nu_N} \left[d_H(\mu, \mu_N^S)^2 \right]^{1/2} \leq C \left\| \mathbb{E}_{\nu_N} [|\Phi - \Phi_N|^2]^{1/2} \right\|_{L_{\mu_0}^2(\mathcal{U})},$$

for general approximations Φ_N . This optimal bound requires that the random misfit Φ_N allows pointwise bounds on $\exp(-\Phi_N)$ and Z_N^S with respect to the distribution ν_N on Φ_N and the Bayesian prior μ_0 on the unknown u . If the distribution of Φ_N does not allow pointwise (L^∞) bounds on $\exp(-\Phi_N)$ and Z_N^S , but only bounds in L^r for some $1 \leq r < \infty$, then the norms in the bounds above need to be strengthened to higher order $L_{\mu_0}^q(\mathcal{U})$ norms and/or higher order moments of the error $\Phi - \Phi_N$, resulting in the quantity $\left\| \mathbb{E}_{\nu_N} [|\Phi - \Phi_N|^p]^{1/p} \right\|_{L_{\mu_0}^q(\mathcal{U})}$ appearing on the right hand sides, for some $q \geq 2$ and $p \geq 1$ (respectively $p \geq 2$).

Our error bounds are explicit in the sense that the aforementioned exponents p and q can typically be calculated explicitly given the structure of Φ_N . This is the case for the GP emulators considered in Stuart and Teckentrup (2017), and also the randomised misfit models and probabilistic numerical solvers

considered here. The constant C in the error bounds, on the other hand, is typically not computable in advance; it involves quantities such as the normalising constants Z and $\mathbb{E}[Z_N^S]$, which for most forward models G are not known analytically and very expensive to compute numerically. In a sense, this is similar to the everyday situation of using an ODE or PDE solver of known order but unknown constant prefactor.

A significant open question in this work is the one highlighted at the end of Section 4: in contrast to randomised dimension reduction using bounded random variables, is the case of *Gaussian* randomly-projected misfits one in which the MAP problem and BIP genuinely have different convergence properties?

Acknowledgements

ALT is partially supported by The Alan Turing Institute under the EPSRC grant EP/N510129/1. HCL and TJS are partially supported by the Freie Universität Berlin within the Excellence Initiative of the German Research Foundation (DFG). This work was partially supported by the DFG through grant CRC 1114 “Scaling Cascades in Complex Systems”, and by the National Science Foundation (NSF) under grant DMS-1127914 to the Statistical and Applied Mathematical Sciences Institute (SAMSI) and SAMSI’s QMC Working Group II “Probabilistic Numerics”. Any opinions, findings, and conclusions or recommendations expressed in this article are those of the authors and do not necessarily reflect the views of the above-named funding agencies and institutions.

References

- C. Andrieu and G. O. Roberts. The pseudo-marginal approach for efficient Monte Carlo computations. *Ann. Statist.*, 37(2):697–725, 2009. doi:10.1214/07-AOS574.
- M. A. Beaumont. Estimation of population growth or decline in genetically monitored populations. *Genetics*, 164(3):1139–1160, 2003. URL <http://www.genetics.org/content/164/3/1139>.
- V. I. Bogachev. *Measure Theory. Vol. I, II*. Springer-Verlag, Berlin, 2007. doi:10.1007/978-3-540-34514-5.
- J. C. Butcher. A stability property of implicit Runge–Kutta methods. *BIT Num. Math.*, 15(4):358–361, Dec 1975. doi:10.1007/BF01931672.
- O. A. Chkrebtii, D. A. Campbell, B. Calderhead, and M. A. Girolami. Bayesian solution uncertainty quantification for differential equations. *Bayesian Anal.*, 11(4):1239–1267, 2016. doi:10.1214/16-BA1017.
- P. R. Conrad, M. Girolami, S. Särkkä, A. M. Stuart, and K. C. Zygalakis. Statistical analysis of differential equations: introducing probability measures on numerical solutions. *Stat. Comput.*, 27(4):1065–1082, 2016. doi:10.1007/s11222-016-9671-0.
- 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(3):424–446, 2013. doi:10.1214/13-STS421.
- M. Dashti and A. M. Stuart. The Bayesian approach to inverse problems. In R. Ghanem, D. Higdon, and H. Owhadi, editors, *Handbook of Uncertainty Quantification*, pages 311–428. Springer, 2016. doi:10.1007/978-3-319-11259-6_7-1.
- M. Dashti, S. Harris, and A. M. Stuart. Besov priors for Bayesian inverse problems. *Inverse Probl. Imaging*, 6(2):183–200, 2012. doi:10.3934/ipi.2012.6.183.
- P. Diaconis. Bayesian numerical analysis. In *Statistical Decision Theory and Related Topics, IV, Vol. 1 (West Lafayette, Ind., 1986)*, pages 163–175. Springer, New York, 1988.
- P. Dupuis and R. S. Ellis. *A Weak Convergence Approach to the Theory of Large Deviations*. Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons, Inc., New York, 1997. doi:10.1002/9781118165904.
- S. N. Evans and P. B. Stark. Inverse problems as statistics. *Inverse Probl.*, 18(4):R55–R97, 2002. doi:10.1088/0266-5611/18/4/201.
- E. Hairer, S. Nørsett, and G. Wanner. *Solving Ordinary Differential Equations I: Nonstiff Problems*, volume 8 of *Springer Series in Computational Mathematics*. Springer-Verlag, New York, 2009. doi:10.1007/978-3-540-78862-1.

- P. Hennig, M. A. Osborne, and M. Girolami. Probabilistic numerics and uncertainty in computations. *Proc. A.*, 471(2179):20150142, 17, 2015. doi:10.1098/rspa.2015.0142.
- D. J. Higham, X. Mao, and A. M. Stuart. Strong convergence of Euler-type methods for nonlinear stochastic differential equations. *SIAM J. Numer. Anal.*, 40(3):1041–1063, 2002. doi:10.1137/S0036142901389530.
- B. Hosseini. Well-posed Bayesian inverse problems with infinitely-divisible and heavy-tailed prior measures. *SIAM/ASA J. Uncertain. Quantif.*, 5(1):1024–1060, 2017. doi:10.1137/16M1096372.
- B. Hosseini and N. Nigam. Well-posed Bayesian inverse problems: priors with exponential tails. *SIAM/ASA J. Uncertain. Quantif.*, 5(1):436–465, 2017. doi:10.1137/16M1076824.
- R. Jordan and D. Kinderlehrer. An extended variational principle. In *Partial Differential Equations and Applications*, volume 177 of *Lecture Notes in Pure and Appl. Math.*, pages 187–200. Dekker, New York, 1996. doi:10.5006/1.3292113.
- J. Kaipio and E. Somersalo. *Statistical and Computational Inverse Problems*, volume 160 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 2005. doi:10.1007/b138659.
- C. H. Kraft. Some conditions for consistency and uniform consistency of statistical procedures. *Univ. California Publ. Statist.*, 2:125–141, 1955.
- M. Lassas and S. Siltanen. Can one use total variation prior for edge-preserving Bayesian inversion? *Inverse Probl.*, 20(5):1537–1563, 2004. doi:10.1088/0266-5611/20/5/013.
- K. Law, A. Stuart, and K. Zygalakis. *Data Assimilation: A Mathematical Introduction*, volume 62 of *Texts in Applied Mathematics*. Springer, 2015. doi:10.1007/978-3-319-20325-6.
- E. B. Le, A. Myers, T. Bui-Thanh, and Q. P. Nguyen. A data-scalable randomized misfit approach for solving large-scale PDE-constrained inverse problems. *Inverse Probl.*, 33(6):065003, 2017. doi:10.1088/1361-6420/aa6cbd.
- H. C. Lie, A. M. Stuart, and T. J. Sullivan. Strong convergence rates of probabilistic integrators for ordinary differential equations, 2017. arXiv:1703.03680.
- A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM J. Optim.*, 19(4):1574–1609, 2008. doi:10.1137/070704277.
- S. Ohta and A. Takatsu. Displacement convexity of generalized relative entropies. *Adv. Math.*, 228(3):1742–1787, 2011. doi:10.1016/j.aim.2011.06.029.
- H. Owhadi and C. Scovel. Qualitative robustness in Bayesian inference. *ESAIM Probab. Stat.*, 21:251–274, 2017. doi:10.1051/ps/2017014.
- H. Owhadi, C. Scovel, and T. J. Sullivan. Brittleness of Bayesian inference under finite information in a continuous world. *Electron. J. Stat.*, 9(1):1–79, 2015. doi:10.1214/15-EJS989.
- M. S. Pinsker. *Information and Information Stability of Random Variables and Processes*. Holden-Day, Inc., San Francisco, Calif.-London-Amsterdam, 1964.
- S. Reich and C. Cotter. *Probabilistic Forecasting and Bayesian Data Assimilation*. Cambridge University Press, New York, 2015. doi:10.1017/CBO9781107706804.
- C. P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer Texts in Statistics. Springer-Verlag, New York, 1999. doi:10.1007/978-1-4757-3071-5.
- M. Schober, D. K. Duvenaud, and P. Hennig. Probabilistic ODE solvers with Runge–Kutta means. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 739–747. Curran Associates, Inc., 2014. https://papers.nips.cc/paper/5451-probabilistic-ode-solvers-with-runge-kutta-means.
- A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on Stochastic Programming: Modeling and Theory*, volume 9 of *MPS/SIAM Series on Optimization*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA; Mathematical Programming Society (MPS), Philadelphia, PA, 2009. doi:10.1137/1.9780898718751.
- J. Skilling. Bayesian solution of ordinary differential equations. In C. R. Smith, G. J. Erickson, and P. O. Neudorfer, editors, *Maximum Entropy and Bayesian Methods*, volume 50 of *Fundamental Theories of Physics*, pages 23–37. Springer, 1992. doi:10.1007/978-94-017-2219-3.
- A. M. Stuart. Inverse problems: a Bayesian perspective. *Acta Numer.*, 19:451–559, 2010. doi:10.1017/S0962492910000061.
- A. M. Stuart and A. L. Teckentrup. Posterior consistency for Gaussian process approximations of Bayesian posterior distributions. *Math. Comput.*, 2017. doi:10.1090/mcom/3244.

Appendix: Proofs of Results

The proofs in this section will make repeated use of the following inequalities for real a and b :

$$(a - b)^2 \leq 2a^2 + 2b^2, \quad (\text{A.1})$$

$$(a - b)^2 = \left(\frac{a^2 - b^2}{a + b} \right)^2 \leq \frac{(a^2 - b^2)^2}{a^2 + b^2}, \quad (\text{A.2})$$

$$|\exp(a) - \exp(b)| \leq (\exp(a) + \exp(b))|a - b|, \quad (\text{A.3})$$

$$[(a + b)ab]^{-1} \leq \max\{a^{-3}, b^{-3}\} \quad \text{for } a, b > 0. \quad (\text{A.4})$$

We also have, for arbitrary $N \in \mathbb{N}$ and $p \geq 1$ (not necessarily integer-valued), by the triangle inequality and Jensen's inequality,

$$\left| \sum_{j=1}^N s_j \right|^p \leq N^p \left(\frac{1}{N} \sum_{j=1}^N |s_j| \right)^p \leq N^{p-1} \sum_{j=1}^N |s_j|^p, \quad (\text{A.5})$$

Proof of Theorem 3.1. Using (2.4) and (3.1), we have

$$\begin{aligned} \sqrt{\frac{d\mu}{d\mu_0}} - \sqrt{\frac{d\mu_N^M}{d\mu_0}} &= \frac{\sqrt{\exp(-\Phi(u))}}{Z^{1/2}} - \frac{\sqrt{\mathbb{E}_{\nu_N}[\exp(-\Phi_N(u))]} }{\mathbb{E}_{\nu_N}[Z_N^S]^{1/2}} \\ &= \frac{\sqrt{\exp(-\Phi(u))} - \sqrt{\mathbb{E}_{\nu_N}[\exp(-\Phi_N(u))]} }{Z^{1/2}} \\ &\quad + \sqrt{\mathbb{E}_{\nu_N}[\exp(-\Phi_N(u))]} \left(\frac{1}{Z^{1/2}} - \frac{1}{\mathbb{E}_{\nu_N}[Z_N^S]^{1/2}} \right). \end{aligned}$$

Inequality (A.1) with $a = Z^{-1/2}(e^{-\Phi(u)/2} - \mathbb{E}_{\nu_N}[e^{-\Phi_N(u)}]^{1/2})$ and $b = \mathbb{E}_{\nu_N}[Z_N^S]^{1/2}(Z^{-1/2} - \mathbb{E}_{\nu_N}[Z_N^S]^{-1/2})$ and the definition (2.1) of the Hellinger distance d_H yield

$$\begin{aligned} 2 d_H(\mu, \mu_N^M)^2 &= \int_{\mathcal{U}} \left(\sqrt{\frac{d\mu}{d\mu_0}}(u) - \sqrt{\frac{d\mu_N^M}{d\mu_0}}(u) \right)^2 d\mu_0(u) \\ &\leq \frac{2}{Z} \left\| \left(\sqrt{\exp(-\Phi)} - \sqrt{\mathbb{E}_{\nu_N}[\exp(-\Phi_N)]} \right)^2 \right\|_{L_{\mu_0}^1} \\ &\quad + 2\mathbb{E}_{\nu_N}[Z_N^S] \left(Z^{-1/2} - \mathbb{E}_{\nu_N}[Z_N^S]^{-1/2} \right)^2 =: I + II. \end{aligned}$$

For the first term, we use inequality (A.2) with $a = e^{-\Phi(u)/2}$ and $b = \mathbb{E}_{\nu_N}[\exp(-\Phi_N(u))]^{1/2}$, together with Hölder's inequality with conjugate exponents p_1 and p'_1 , to derive

$$\begin{aligned} \frac{Z}{2} I &\leq \left\| (\exp(-\Phi) - \mathbb{E}_{\nu_N}[\exp(-\Phi_N)])^2 (\exp(-\Phi) + \mathbb{E}_{\nu_N}[\exp(-\Phi_N)])^{-1} \right\|_{L_{\mu_0}^1} \\ &\leq \left\| (\exp(-\Phi) - \mathbb{E}_{\nu_N}[\exp(-\Phi_N)])^2 \right\|_{L_{\mu_0}^{p'_1}} \left\| (\exp(-\Phi) + \mathbb{E}_{\nu_N}[\exp(-\Phi_N)])^{-1} \right\|_{L_{\mu_0}^{p_1}}. \quad (\text{A.6}) \end{aligned}$$

We estimate the second factor on the right-hand side of (A.6). Using the facts that $x \mapsto 1/x$ is decreasing on $(0, \infty)$, that $(x + y)^{-1} \leq \min\{x^{-1}, y^{-1}\}$ for all $x, y > 0$, and that both $\exp(-\Phi(u))$ and $\mathbb{E}_{\nu_N}[\exp(-\Phi_N(u))]$ are strictly positive, we obtain

$$\left\| (\exp(-\Phi) + \mathbb{E}_{\nu_N}[\exp(-\Phi_N)])^{-1} \right\|_{L_{\mu_0}^{p_1}} \leq \left\| \min\{\exp(\Phi), \mathbb{E}_{\nu_N}[\exp(-\Phi_N)]^{-1}\} \right\|_{L_{\mu_0}^{p_1}}.$$

For $f, g \in L_{\mu_0}^1(\mathcal{U})$, the partition $\mathcal{U} = \{f < g\} \uplus \{f \geq g\}$ and the corresponding integral inequalities on $\{f < g\}$ and $\{f \geq g\}$ imply that $\|\min\{f, g\}\|_{L_{\mu_0}^1} \leq \min\{\|f\|_{L_{\mu_0}^1}, \|g\|_{L_{\mu_0}^1}\}$. Hence,

$$\left\| \min\left\{e^{-\Phi}, \mathbb{E}_{\nu_N}[e^{-\Phi_N}]^{-1}\right\} \right\|_{L_{\mu_0}^{p_1}} \leq \min\left\{\|e^{\Phi}\|_{L_{\mu_0}^{p_1}}, \left\|\mathbb{E}_{\nu_N}[e^{-\Phi_N}]^{-1}\right\|_{L_{\mu_0}^{p_1}}\right\} \leq C_1, \quad (\text{A.7})$$

where $C_1 = C_1(p_1)$ is the constant specified in assumption (a). This completes our estimate for the second factor on the right-hand side of (A.6). For the first factor, the linearity of expectation, inequality (A.3), and Hölder's inequality with conjugate exponents p_2, p_2' with respect to ν_N and p_3, p_3' with respect to μ_0 give

$$\begin{aligned} & \left\| (\exp(-\Phi) - \mathbb{E}_{\nu_N}[\exp(-\Phi_N)])^2 \right\|_{L_{\mu_0}^{p_1'}} = \left\| \mathbb{E}_{\nu_N}[\exp(-\Phi) - \exp(-\Phi_N)]^2 \right\|_{L_{\mu_0}^{p_1'}} \\ & \leq \left\| \mathbb{E}_{\nu_N}[|\exp(-\Phi) + \exp(-\Phi_N)| |\Phi - \Phi_N|] \right\|_{L_{\mu_0}^{p_1'}}^2 \\ & \leq \left\| \mathbb{E}_{\nu_N}[(\exp(-\Phi) + \exp(-\Phi_N))^{p_2}]^{2/p_2} \mathbb{E}_{\nu_N}[|\Phi - \Phi_N|^{p_2'}]^{2/p_2'} \right\|_{L_{\mu_0}^{p_1'}} \\ & \leq \left\| \mathbb{E}_{\nu_N}[(\exp(-\Phi) + \exp(-\Phi_N))^{p_2}]^{1/p_2} \right\|_{L_{\mu_0}^{2p_1'p_3}}^2 \left\| \mathbb{E}_{\nu_N}[|\Phi - \Phi_N|^{p_2'}]^{1/p_2'} \right\|_{L_{\mu_0}^{2p_1'p_3}}^2 \end{aligned} \quad (\text{A.8})$$

Letting $C_2 = C_2(p_1', p_2, p_3)$ be the constant in assumption (b), and using (A.7), it follows that

$$I \leq \frac{2}{Z} \cdot C_1(p_1) \cdot C_2^2(p_1', p_2, p_3) \cdot \left\| \mathbb{E}_{\nu_N}[|\Phi - \Phi_N|^{p_2'}]^{1/p_2'} \right\|_{L_{\mu_0}^{2p_1'p_3}}^2.$$

Now inequality (A.2) with $a = \mathbb{E}_{\nu_N}[Z_N^S]^{-1/2}$ and $b = Z^{-1/2}$ and inequality (A.4) yield

$$\begin{aligned} \frac{1}{2\mathbb{E}_{\nu_N}[Z_N^S]} II &= \left(Z^{-1/2} - (\mathbb{E}_{\nu_N}[Z_N^S])^{-1/2} \right)^2 \\ &= \left(\frac{\mathbb{E}_{\nu_N}[Z_N^S] - Z}{Z\mathbb{E}_{\nu_N}[Z_N^S]} \right)^2 \frac{Z\mathbb{E}_{\nu_N}[Z_N^S]}{Z + \mathbb{E}_{\nu_N}[Z_N^S]} \\ &\leq (\mathbb{E}_{\nu_N}[Z_N^S] - Z)^2 \max\{Z^{-3}, \mathbb{E}_{\nu_N}[Z_N^S]^{-3}\} \\ &\leq (\mathbb{E}_{\nu_N}[Z_N^S] - Z)^2 \max\{Z^{-3}, C_3^{-3}\}, \end{aligned}$$

where the last inequality follows from assumption (c).

Using Tonelli's theorem, Jensen's inequality, inequality (A.3), and Hölder's inequality with the same conjugate exponent pairs that we used to obtain (A.8),

$$\begin{aligned} & (\mathbb{E}_{\nu_N}[Z_N^S] - Z)^2 \\ &= \mathbb{E}_{\mu_0}[\mathbb{E}_{\nu_N}[\exp(-\Phi_N) - \exp(-\Phi)]]^{2p_1'/p_1'} \\ &\leq \left\| \mathbb{E}_{\nu_N}[\exp(-\Phi) - \exp(-\Phi_N)] \right\|_{L_{\mu_0}^{p_1'}}^2 \\ &\leq \left\| \mathbb{E}_{\nu_N}[(\exp(-\Phi) + \exp(-\Phi_N))^{p_2}]^{1/p_2} \right\|_{L_{\mu_0}^{2p_1'p_3}}^2 \left\| \mathbb{E}_{\nu_N}[|\Phi - \Phi_N|^{p_2'}]^{1/p_2'} \right\|_{L_{\mu_0}^{2p_1'p_3}}^2 \\ &\leq C_2^2(p_1, p_2, p_3) \left\| \mathbb{E}_{\nu_N}[|\Phi - \Phi_N|^{p_2'}]^{1/p_2'} \right\|_{L_{\mu_0}^{2p_1'p_3}}^2, \end{aligned}$$

where assumption (b) yields the last inequality. Combining the estimates for I and II yields (3.3). $\square$

Proof of Theorem 3.2. This proof is similar to the proof of Theorem 3.1. Since

$$\sqrt{\frac{d\mu}{d\mu_0}} - \sqrt{\frac{d\mu_N^S}{d\mu_0}} = \frac{e^{-\Phi(u)/2} - e^{-\Phi_N(u)/2}}{Z^{1/2}} - e^{-\Phi_N(u)/2} \left(\frac{1}{\sqrt{Z_N^S}} - \frac{1}{Z^{1/2}} \right),$$

Tonelli's theorem, inequality (A.1), and Jensen's inequality yield

$$\begin{aligned}
\mathbb{E}_{\nu_N}[d_H(\mu, \mu_N^S)^2] &= \frac{1}{2} \left\| \mathbb{E}_{\nu_N} \left[\left(\sqrt{\frac{d\mu}{d\mu_0}} - \sqrt{\frac{d\mu_N^S}{d\mu_0}} \right)^2 \right] \right\|_{L_{\mu_0}^1} \\
&\leq \frac{1}{Z} \left\| \mathbb{E}_{\nu_N} \left[\left(\sqrt{\exp(-\Phi)} - \sqrt{\exp(-\Phi_N)} \right)^2 \right] \right\|_{L_{\mu_0}^1} \\
&\quad + \mathbb{E}_{\nu_N} \left[Z_N^S (Z^{-1/2} - (Z_N^S)^{-1/2})^2 \right] \\
&=: I + II.
\end{aligned}$$

For the first term I , inequality (A.3), and Hölder's inequality with conjugate exponent pairs (q_1, q'_1) and (q_2, q'_2) give

$$\begin{aligned}
ZI &= \left\| \mathbb{E}_{\nu_N} \left[\left(\sqrt{\exp(-\Phi)} - \sqrt{\exp(-\Phi_N)} \right)^2 \right] \right\|_{L_{\mu_0}^1} \\
&\leq \frac{1}{4} \left\| \mathbb{E}_{\nu_N} [|\exp(-\Phi/2) + \exp(-\Phi_N/2)|^2 |\Phi - \Phi_N|^2] \right\|_{L_{\mu_0}^1} \\
&\leq \left\| \mathbb{E}_{\nu_N} [|\exp(-\Phi/2) + \exp(-\Phi_N/2)|^{2q_1}]^{1/q_1} \mathbb{E}_{\nu_N} [|\Phi - \Phi_N|^{2q'_1}]^{1/q'_1} \right\|_{L_{\mu_0}^1} \\
&\leq \left\| \mathbb{E}_{\nu_N} [(\exp(-\Phi/2) + \exp(-\Phi_N/2))^{2q_1}]^{1/q_1} \right\|_{L_{\mu_0}^{q_2}} \left\| \mathbb{E}_{\nu_N} [|\Phi - \Phi_N|^{2q'_1}]^{1/2q'_1} \right\|_{L_{\mu_0}^{2q'_2}}^2.
\end{aligned}$$

By (a), we may bound the first factor on the right-hand side of the last inequality by $D_1(q_1, q_2)$. Now by (A.2) with $a = Z^{-1/2}$ and $b = (Z_N^S)^{-1/2}$, and by inequality (A.4), we obtain (see the proof of Theorem 3.1 after (A.8)) that

$$II \leq \mathbb{E}_{\nu_N} \left[Z_N^S \max\{Z^{-3}, (Z_N^S)^{-3}\} (Z - Z_N^S)^2 \right].$$

Jensen's inequality and another application of inequality (A.2) yield

$$(Z - Z_N^S)^2 \leq \|\exp(-\Phi) - \exp(-\Phi_N)\|_{L_{\mu_0}^2}^2 \leq \|(\exp(-\Phi) + \exp(-\Phi_N))^2 (\Phi - \Phi_N)^2\|_{L_{\mu_0}^1}.$$

Combining the preceding two estimates, using Tonelli's theorem and Hölder's inequality with the same conjugate exponent pairs (q_1, q'_1) and (q_2, q'_2) as used in the bound for I , and using (b), we get

$$\begin{aligned}
II &\leq \left\| \mathbb{E}_{\nu_N} \left[Z_N^S \max\{Z^{-3}, (Z_N^S)^{-3}\} (e^{-\Phi} + e^{-\Phi_N})^2 (\Phi - \Phi_N)^2 \right] \right\|_{L_{\mu_0}^1} \\
&\leq \left\| \mathbb{E}_{\nu_N} \left[\left(Z_N^S \max\{Z^{-3}, (Z_N^S)^{-3}\} (e^{-\Phi} + e^{-\Phi_N})^2 \right)^{q_1} \right]^{\frac{1}{q_1}} \mathbb{E}_{\nu_N} [|\Phi - \Phi_N|^{2q'_1}]^{\frac{1}{q'_1}} \right\|_{L_{\mu_0}^1} \\
&\leq D_2(q_1, q_2) \left\| \mathbb{E}_{\nu_N} [|\Phi - \Phi_N|^{2q'_1}]^{1/2q'_1} \right\|_{L_{\mu_0}^{2q'_2}}^2.
\end{aligned}$$

Combining the preceding estimates yields (3.4). $\square$

Proof of Lemma 3.4. Since $\exp(\Phi) \in L_{\mu_0}^{p^*}$, examination of assumption (a) of Theorem 3.1 indicates that we may set $p_1 = p^*$ and $C_1 := \|\exp(\Phi)\|_{L_{\mu_0}^{p^*}}$. By (3.5), it follows that $\mathbb{E}_{\nu_N}[\exp(-\Phi) + \exp(-\Phi_N)] \leq 2\exp(C_0)$; thus assumption (b) of Theorem 3.1 holds with $p_2 = p_3 = +\infty$ (so that $2p'_1 p_3 = +\infty$) and $C_2 = 2\exp(C_0)$. We now prove that Assumption (c) of Theorem 3.1 holds. It follows by setting $x = -\Phi$ and $y = -\Phi_N$ in inequality (A.3) that $|\exp(-\Phi) - \exp(-\Phi_N)| \leq 2\exp(C_0)|\Phi - \Phi_N|$. Thus

$$\begin{aligned}
|Z_N^S - Z| &= |\mathbb{E}_{\mu_0}[\exp(-\Phi_N) - \exp(-\Phi)]| \\
&\leq \mathbb{E}_{\mu_0}[|\exp(-\Phi_N) - \exp(-\Phi)|] \\
&\leq 2\exp(C_0)\mathbb{E}_{\mu_0}[|\Phi - \Phi_N|].
\end{aligned} \tag{A.9}$$

Using Jensen's inequality, (A.9), Tonelli's theorem, and (3.6),

$$|\mathbb{E}_{\nu_N}[Z_N^S] - Z| \leq \mathbb{E}_{\nu_N}[|Z_N^S - Z|] \leq 2e^{C_0} \|\mathbb{E}_{\nu_N}[|\Phi - \Phi_N|]\|_{L_{\mu_0}^1} \leq \min\left\{Z - \frac{1}{C_3}, C_3 - Z\right\}.$$

The last inequality implies that assumption (c) of Theorem 3.1 holds with the same C_3 as in (3.6), since for any $0 < C_3 < +\infty$ that satisfies $C_3^{-1} < Z < C_3$ and (3.6), we have

$$C_3^{-1} - Z \leq \mathbb{E}_{\nu_N}[Z_N^S] - Z \leq Z - C_3^{-1} \implies C_3^{-1} \leq \mathbb{E}_{\nu_N}[Z_N^S]$$

and

$$Z - C_3 \leq \mathbb{E}_{\nu_N}[Z_N^S] - Z \leq C_3 - Z \implies \mathbb{E}_{\nu_N}[Z_N^S] \leq C_3,$$

and combining both the implied statements yields assumption (c) of Theorem 3.1; thus (3.3) holds, as desired.

Now note that (3.5) implies that assumption (a) of Theorem 3.2 holds with $q_1 = q_2 = +\infty$ and $D_1 = 4\exp(C_0)$. Furthermore, (3.5) also implies that $Z_N^S = \mathbb{E}_{\mu_0}[\exp(-\Phi_N)] \leq \exp(C_0)$ for all Φ_N . Thus, given that Z is ν_N -a.s. constant, and given that there exists some $0 < C_3 < \infty$ such that $C_3^{-1} < Z < C_3$,

$$\begin{aligned} & \mathbb{E}_{\nu_N}\left[(Z_N^S)^{q_1} \max\{Z^{-3}, (Z_N^S)^{-3}\}^{q_1} (\exp(-\Phi(u)) + \exp(-\Phi_N(u)))^{2q_1}\right]^{1/q_1} \\ & \leq 4\exp(3C_0)\mathbb{E}_{\nu_N}\left[\max\{C_3^{-3}, (Z_N^S)^{-3}\}^{q_1}\right]^{1/q_1}. \end{aligned} \quad (\text{A.10})$$

A necessary and sufficient condition for setting $q_1 = +\infty$ above (and therefore also in assumption (b) of Theorem 3.2) is that Z_N^S is ν_N -a.s. bounded away from zero by a constant that does not depend on N . By the convexity and monotonicity of $x \mapsto \exp(x)$,

$$Z_N^S = \mathbb{E}_{\mu_0}[\exp(-\Phi_N)] \geq \exp(\mathbb{E}_{\mu_0}[-\Phi_N]) \geq \exp(-C_4),$$

for C_4 as in (3.7). In particular, if (3.7) holds, then so does assumption (b) of Theorem 3.2, with $q_1 = q_2 = +\infty$ and $D_2 = 4\exp(3C_0)\max\{C_3^{-3}, \exp(3C_4)\}$, by inequality (A.10). $\square$

Proof of Lemma 3.5. The proof proceeds in the same way as the proof of Lemma 3.4, with the exception that we need to prove that the assumption that $\mathbb{E}_{\nu_N}[\exp(\rho^*\Phi_N)] \in L_{\mu_0}^1$ for some $\rho^* > 2$ implies that assumption (a) of Theorem 3.1 and assumption (b) of Theorem 3.2 hold with the stated parameters. Therefore, the proof will only concern these two assertions. Since $x \mapsto x^{-t}$ is strictly convex on $\mathbb{R}_{>0}$ for any $t > 0$, Jensen's inequality yields that $\|\mathbb{E}_{\nu_N}[\exp(-\Phi_N)]^{-1}\|_{L_{\mu_0}^t} \leq \|\mathbb{E}_{\nu_N}[\exp(t\Phi_N)]\|_{L_{\mu_0}^1}^{1/t}$. Therefore, setting $t = \rho^*$, we find that assumption (a) of Theorem 3.1 holds, with $p_1 = \rho^*$ and $C_1 = \|\mathbb{E}_{\nu_N}[\exp(\rho^*\Phi_N)]\|_{L_{\mu_0}^1}^{1/\rho^*}$. The inequality $\max\{x, y\} \leq x + y$ for $x, y \geq 0$ implies that

$$\begin{aligned} & \mathbb{E}_{\nu_N}\left[\max\{Z_N^S Z^{-3}, (Z_N^S)^{-2}\}^{q_1} (\exp(-\Phi(u)) + \exp(-\Phi_N(u)))^{2q_1}\right]^{1/q_1} \\ & \leq 4\exp(2C_0)\left(C_3^{-3}\exp(C_0) + \mathbb{E}_{\nu_N}\left[(Z_N^S)^{-2q_1}\right]^{1/q_1}\right), \end{aligned}$$

while Jensen's inequality, Tonelli's theorem, and the definition of the $L_{\mu_0}^1$ -norm yield that

$$\begin{aligned} \mathbb{E}_{\nu_N}[(Z_N^S)^{-2q_1}] & \leq \mathbb{E}_{\nu_N}[\mathbb{E}_{\mu_0}[\exp(2q_1\Phi_N)]] \\ & = \mathbb{E}_{\mu_0}[\mathbb{E}_{\nu_N}[\exp(2q_1\Phi_N)]] \\ & = \|\mathbb{E}_{\nu_N}[\exp(2q_1\Phi_N)]\|_{L_{\mu_0}^1}. \end{aligned}$$

Since the last term is finite for $q_1 \leq \rho^*/2$ by the hypothesis that $\mathbb{E}_{\nu_N}[\exp(\rho^*\Phi_N)] \in L_{\mu_0}^1$, it follows that assumption (b) of Theorem 3.2 holds with the parameters $q_1 = \rho^*/2$, $q_2 = +\infty$, and the scalar $D_2 = 4\exp(2C_0)(C_3^{-3}\exp(C_0) + \|\mathbb{E}_{\nu_N}[\exp(\rho^*\Phi_N)]\|_{L_{\mu_0}^1}^{2/\rho^*})$. $\square$

Proof of Proposition 3.6. Recall (3.8), and fix an arbitrary $u \in \mathcal{U}$. We have

$$|\Phi(u) - \Phi_N(u)| = \frac{1}{2} |\langle G(u) - y, \Gamma^{-1}(G(u) - y) \rangle - \langle G_N(u) - y, \Gamma^{-1}(G_N(u) - y) \rangle|.$$

Adding and subtracting $\langle G_N(u) - y, \Gamma^{-1}(G(u) - y) \rangle$ inside the absolute value, rearranging terms, applying the Cauchy-Schwarz inequality, and letting C_Γ be the largest eigenvalue of Γ^{-1} yields

$$\begin{aligned} |\Phi(u) - \Phi_N(u)| &= \frac{1}{2} |\langle \Gamma^{-1}(G(u) - y), G(u) - G_N(u) \rangle + \langle \Gamma^{-1}(G_N(u) - y), G(u) - G_N(u) \rangle| \\ &= \frac{1}{2} |\langle G(u) - y + G_N(u) - y, \Gamma^{-1}(G(u) - G_N(u)) \rangle| \\ &\leq C_\Gamma \|G(u) + G_N(u) - 2y\| \|G(u) - G_N(u)\|. \end{aligned} \quad (\text{A.11})$$

By the triangle inequality,

$$\|G(u) + G_N(u) - 2y\| \leq 2 \max\{\|G(u) - y\|, \|G_N(u) - y\|\} = 2 \max\{\Phi(u)^{1/2}, \Phi_N(u)^{1/2}\},$$

and the triangle inequality and (3.8) yield

$$\begin{aligned} \Phi_N(u)^{1/2} &= 2^{-1/2} \|G_N(u) - y\| \\ &= 2^{-1/2} \|G(u) - y + G_N(u) - G(u)\| \\ &\leq 2^{-1/2} (2^{1/2} \Phi(u)^{1/2} + \|G_N(u) - G(u)\|) \\ &= \Phi(u)^{1/2} + 2^{-1/2} \|G_N(u) - G(u)\|. \end{aligned}$$

Together, these inequalities yield

$$\|G(u) - y + G_N(u) - y\| \leq 2(\Phi(u)^{1/2} + 2^{-1/2} \|G_N(u) - G(u)\|),$$

and substituting the above into (A.11) yields

$$|\Phi(u) - \Phi_N(u)| \leq 2C_\Gamma \left(\Phi(u)^{1/2} \|G_N(u) - G(u)\| + \|G(u) - G_N(u)\|^2 \right),$$

thus proving (3.9). Using (A.5) yields

$$|\Phi(u) - \Phi_N(u)|^q \leq 2^{q-1} (2C_\Gamma)^q \left(\Phi(u)^{q/2} \|G_N(u) - G(u)\|^q + \|G(u) - G_N(u)\|^{2q} \right).$$

Now take expectations with respect to ν_N : since G and Φ are constant with respect to ν_N ,

$$\mathbb{E}_{\nu_N} [|\Phi(u) - \Phi_N(u)|^q] \leq (4C_\Gamma)^q \left(\Phi(u)^{q/2} \mathbb{E}_{\nu_N} [\|G_N(u) - G(u)\|^q] + \mathbb{E}_{\nu_N} [\|G(u) - G_N(u)\|^{2q}] \right),$$

and taking the q^{th} root of both sides proves (3.10). $\square$

Proof of Corollary 3.7. Taking the $L_{\mu_0}^s$ norm of both sides of the second inequality in Proposition 3.6, and applying (A.5) with $s/q \geq 1$, we obtain

$$\begin{aligned} &\left\| \mathbb{E}_{\nu_N} [|\Phi - \Phi_N|^q]^{1/q} \right\|_{L_{\mu_0}^s} \\ &\leq (4C_\Gamma) \mathbb{E}_{\mu_0} \left[\left(\Phi^{q/2} \mathbb{E}_{\nu_N} [\|G_N - G\|^q] + \mathbb{E}_{\nu_N} [\|G_N - G\|^{2q}] \right)^{s/q} \right]^{1/s} \\ &\leq (4C_\Gamma) 2^{1/q-1/s} \left(\mathbb{E}_{\mu_0} \left[\Phi(u)^{s/2} \mathbb{E}_{\nu_N} [\|G_N - G\|^{s/q}] \right] + \mathbb{E}_{\mu_0} \left[\mathbb{E}_{\nu_N} [\|G_N - G\|^{2q}]^{s/q} \right] \right)^{1/s}. \end{aligned}$$

By the Cauchy-Schwarz inequality and Jensen's inequality,

$$\begin{aligned} \mathbb{E}_{\mu_0} \left[\Phi^{s/2} \mathbb{E}_{\nu_N} [\|G_N - G\|^{s/q}] \right] &\leq \left(\mathbb{E}_{\mu_0} [\Phi^s] \mathbb{E}_{\mu_0} \left[\mathbb{E}_{\nu_N} [\|G_N - G\|^{2s/q}] \right] \right)^{1/2} \\ &\leq \left(\mathbb{E}_{\mu_0} [\Phi^s] \mathbb{E}_{\mu_0} \left[\mathbb{E}_{\nu_N} [\|G_N - G\|^{2q}]^{s/q} \right] \right)^{1/2}. \end{aligned}$$

Since $0 \leq a \leq 1 \implies a \leq a^{1/2}$, the hypotheses of the corollary and the preceding imply that

$$\|\mathbb{E}_{\nu_N} [|\Phi - \Phi_N|^q]^{1/q}\|_{L_{\mu_0}^s} \leq (4C_\Gamma) 2^{1/q-1/s} \left(\mathbb{E}_{\mu_0} [\Phi^s]^{1/2} + 1 \right)^{1/s} \|\mathbb{E}_{\nu_N} [\|G_N - G\|^{2q}]^{1/q}\|_{L_{\mu_0}^s}^{1/2}.$$

Since $2^{1/q-1/s} \leq 2^{1/q} \leq 2$, the proof is complete. $\square$

Proof of Lemma 3.8. Given (3.8), we may choose the parameter C_0 in (3.5) to be $C_0 = 0$. By Jensen's inequality, (3.11) implies (3.6). $\square$

Proof of Theorem 3.9. We first verify that Assumption 3.3 holds. Since Φ and Φ_N satisfy (3.8), it follows that we may set $C_0 = 0$ in (3.5). Since we assume throughout that $0 < Z = \mathbb{E}_{\mu_0} [\exp(-\Phi)] < \infty$, it follows that Φ has moments of all orders, and hence belongs to $L_{\mu_0}^s$ for all $s \in \mathbb{N}$. Therefore, given that (3.11) holds for $q, s \geq 1$, it follows from Jensen's inequality and Corollary 3.7 that we can make $\|\mathbb{E}_{\nu_N} [\Phi_N - \Phi]\|_{L_{\mu_0}^1}$ as small as desired. In particular, for any $0 < C_3 < +\infty$ that satisfies $C_3^{-1} < Z < C_3$, there exists a $N^*(C_3) \in \mathbb{N}$ such that, for all $N \geq N^*(C_3)$, (3.6) holds.

The rest of the proof consists of applying Lemma 3.4 or 3.5, Corollary 3.7 and Lemma 3.8.

Case (a). The hypotheses in this case ensure that we may apply Lemma 3.4. Set $p_1 = p^*$ and $p_2 = p_3 = +\infty$, so that $p'_1 = (p^*)' = p^*/(p^* - 1)$ and $p'_2 = p'_3 = 1$. Substituting these exponents into (3.3a) and applying Corollary 3.7 with $s = 2p'_1 p'_3 = 2p^*/(p^* - 1)$ and $q = p'_2 = 1$ (note that $s \geq q \geq 1$), we obtain

$$d_H(\mu, \mu_N^M) \leq C \|\mathbb{E}_{\nu_N} [|\Phi - \Phi_N|]\|_{L_{\mu_0}^{2p^*/(p^*-1)}(\mathcal{U})} \leq C \|\mathbb{E}_{\nu_N} [\|G - G_N\|^2]\|_{L_{\mu_0}^{2p^*/(p^*-1)}(\mathcal{U})}^{1/2},$$

where $C > 0$ changes value between inequalities. Thus we have shown that (3.12) holds with $r_1 = 1$ and $r_2 = 2p^*/(p^* - 1)$.

To prove that (3.13) holds with the desired exponents, we again use Lemma 3.4 to set $q_1 = q_2 = +\infty$, so that $q'_1 = q'_2 = 1$. Substituting these exponents into (3.4), and applying Corollary 3.7 with $s = 2q'_2 = 2$ and $q = 2q'_1 = 2$, we obtain

$$\mathbb{E}_{\nu_N} [d_H(\mu, \mu_N^S)^2]^{1/2} \leq D \|\mathbb{E}_{\nu_N} [|\Phi - \Phi_N|^2]\|_{L_{\mu_0}^2}^{1/2} \leq D \|\mathbb{E}_{\nu_N} [\|G - G_N\|^4]\|_{L_{\mu_0}^2(\mathcal{U})}^{1/2},$$

where $D > 0$ changes value between inequalities. Thus we have shown that (3.13) holds with $s_1 = s_2 = 2$.

It remains to ensure that both the rightmost terms above converge to zero. Since (3.11) holds with $q = 2$ and $s = 2p^*/(p^* - 1)$, the desired convergence follows from the nesting property of finite-measure L^p -spaces. Therefore, both μ_N^M and μ_N^S converge to μ as claimed.

Case (b). Since the arguments in this case are the same as in the previous case, we only record the different material.

The hypotheses ensure that we may apply Lemma 3.5. Set $p_1 = \rho^*$ and $p_2 = p_3 = +\infty$, so that $(p_1)' = \rho^*/(\rho^* - 1)$ and $p'_2 = p'_3 = 1$. Substituting these exponents into (3.3a) and applying Corollary 3.7 with $s = 2p'_1 p'_3 = 2\rho^*/(\rho^* - 1)$ and $q = p'_2 = 1$, we obtain

$$d_H(\mu, \mu_N^M) \leq C \|\mathbb{E}_{\nu_N} [|\Phi - \Phi_N|]\|_{L_{\mu_0}^{2\rho^*/(\rho^*-1)}(\mathcal{U})} \leq C \|\mathbb{E}_{\nu_N} [\|G - G_N\|^2]\|_{L_{\mu_0}^{2\rho^*/(\rho^*-1)}(\mathcal{U})}^{1/2},$$

where $C > 0$ changes value between inequalities. Thus we have shown that (3.12) holds with $r_1 = 1$ and $r_2 = 2\rho^*/(\rho^* - 1)$.

To prove that (3.13) holds with the desired exponents, we again use Lemma 3.5 to set $q_1 = \frac{\rho^*}{2}$ and $q_2 = +\infty$, so that $q'_1 = \rho^*/(\rho^* - 2)$ and $q'_2 = 1$. Substituting these exponents into (3.4), and applying Corollary 3.7 with $s = 2q'_2 = 2$ and $q = 2q'_1 = 2\rho^*/(\rho^* - 2)$, we obtain

$$\begin{aligned} \mathbb{E}_{\nu_N} [d_H(\mu, \mu_N^S)^2]^{1/2} &\leq D \|\mathbb{E}_{\nu_N} [|\Phi - \Phi_N|^{2\rho^*/(\rho^*-2)}]^{(\rho^*-2)/(2\rho^*)}\|_{L_{\mu_0}^2} \\ &\leq D \|\mathbb{E}_{\nu_N} [\|G - G_N\|^{4\rho^*/(\rho^*-2)}]^{(\rho^*-2)/(2\rho^*)}\|_{L_{\mu_0}^2(\mathcal{U})}^{1/2}, \end{aligned}$$

where $D > 0$ changes value between inequalities. Thus (3.13) holds with $s_1 = 2\rho^*/(\rho^* - 2)$ and $s_2 = 2$. Since (3.11) holds with $q = 2\rho^*/(\rho^* - 2)$ and $s = 2\rho^*/(\rho^* - 1)$, it follows from the nesting property of L^p -spaces defined on finite measure spaces that both

$$\left\| \mathbb{E}_{\nu_N} [\|G - G_N\|^2] \right\|_{L_{\mu_0}^{2\rho^*/(\rho^*-1)}(\mathcal{U})}^{1/2} \text{ and } \left\| \mathbb{E}_{\nu_N} [\|G - G_N\|^{4\rho^*/(\rho^*-2)}] \right\|_{L_{\mu_0}^2(\mathcal{U})}^{(\rho^*-2)/(2\rho^*)} \right\|^{1/2}$$

converge to zero. $\square$

Proof of Proposition 4.1. We start by verifying the assumptions of Theorem 3.2. First, since $\Phi(u) \geq 0$ for all $u \in \mathcal{U}$, and $\Phi_N(u) \geq 0$ for all $u \in \mathcal{U}$ and all $\{\sigma^{(i)}\}_{i=1}^N$, assumption (a) is satisfied for $q_1 = q_2 = \infty$. For assumption (b), we then have, for any $q_2 \in [1, \infty]$,

$$\begin{aligned} & \left\| \left(\mathbb{E}_\sigma \left[(Z_N^S)^{q_1} \max\{Z^{-3}, (Z_N^S)^{-3}\}^{q_1} (\exp(-\Phi(u)) + \exp(-\Phi_N(u)))^{2q_1} \right] \right)^{1/q_1} \right\|_{L_{\mu_0}^{q_2}(\mathcal{U})} \\ & \leq 4 \mathbb{E}_\sigma \left[(Z_N^S)^{q_1} \max\{Z^{-3}, (Z_N^S)^{-3}\}^{q_1} \right]^{1/q_1} \\ & \leq 4 \left(Z^{-3q_1} \mathbb{E}_\sigma [(Z_N^S)^{q_1}] + \mathbb{E}_\sigma [(Z_N^S)^{-2q_1}] \right)^{1/q_1}. \end{aligned}$$

Since $\Phi_N(u) \geq 0$ for all $u \in \mathcal{U}$ and all $\{\sigma^{(i)}\}_{i=1}^N$, we have for any $q_1 \in [1, \infty]$

$$\mathbb{E}_\sigma [(Z_N^S)^{q_1}]^{1/q_1} = \mathbb{E}_\sigma \left[\left(\int_{\mathcal{U}} \exp(-\Phi_N(u)) d\mu_0(u) \right)^{q_1} \right]^{1/q_1} \leq 1. \quad (\text{A.12})$$

Using the ℓ -sparse distribution of σ , we further have $|\sigma_j^{(i)}| \leq \sqrt{s}$ and

$$\Phi_N(u) = \frac{1}{2N} \sum_{i=1}^N |\sigma^{(i)\top} (\Gamma^{-1/2}(y - G(u)))|^2 \leq \frac{s}{2} \|\Gamma^{-1/2}(y - G(u))\|^2 = s\Phi(u),$$

which implies that $Z_N^S \geq Z_s = \int_{\mathcal{U}} \exp(-s\Phi(u)) d\mu_0(u)$. It follows that, for any $q_1 \in [1, \infty]$,

$$\mathbb{E}_\sigma [(Z_N^S)^{-2q_1}]^{1/q_1} \leq \mathbb{E}_\sigma [Z_s^{-2q_1}]^{1/q_1} = Z_s^{-2},$$

and assumption (b) is hence also satisfied for $q_1 = q_2 = \infty$. Hence, by Theorem 3.2,

$$\left(\mathbb{E}_\sigma [d_H(\mu, \mu_N^S)^2] \right)^{1/2} \leq C \left\| \left(\mathbb{E}_\sigma [|\Phi(u) - \Phi_N(u)|^2] \right)^{1/2} \right\|_{L_{\mu_0}^2(\mathcal{U})}.$$

Using standard properties of Monte Carlo estimators (see e.g. Robert and Casella (1999)), we have

$$\left(\mathbb{E}_\sigma [|\Phi(u) - \Phi_N(u)|^2] \right)^{1/2} = \sqrt{\frac{\mathbb{V}_\sigma \left[\frac{1}{2} |\sigma^\top \Gamma^{-1/2}(y - G(u))|^2 \right]}{N}}.$$

Now, using $\mathbb{V}[X] = \mathbb{E}[X^2] - \mathbb{E}[X]^2$, $\left(\sum_{j=1}^J x_j \right)^4 \leq J^3 \sum_{j=1}^J x_j^4$, the linearity of expectation, the ℓ -sparse

distribution of σ , and $\|x\|_4 \leq \|x\|_2$, we have

$$\begin{aligned}
0 &\leq \mathbb{V}_\sigma \left[\frac{1}{2} |\sigma^\top \Gamma^{-1/2} (y - G(u))|^2 \right] \\
&= \mathbb{E}_\sigma \left[\frac{1}{4} |\sigma^\top \Gamma^{-1/2} (y - G(u))|^4 \right] - \mathbb{E}_\sigma \left[\frac{1}{4} |\sigma^\top \Gamma^{-1/2} (y - G(u))|^2 \right]^2 \\
&= \mathbb{E}_\sigma \left[\frac{1}{4} \left| \sum_{j=1}^J \sigma_j (\Gamma^{-1/2} (y - G(u)))_j \right|^4 \right] - \frac{1}{4} \|\Gamma^{-1/2} (y - G(u))\|^4 \\
&\leq \frac{1}{4} J^3 \sum_{j=1}^J \mathbb{E}_\sigma [\sigma_j^4] (\Gamma^{-1/2} (y - G(u)))_j^4 - \frac{1}{4} \|\Gamma^{-1/2} (y - G(u))\|^4 \\
&= \frac{1}{4} J^3 \mathbb{E}_\sigma [\sigma_j^4] \|\Gamma^{-1/2} (y - G(u))\|_4^4 - \frac{1}{4} \|\Gamma^{-1/2} (y - G(u))\|^4 \\
&\leq (J^3 \mathbb{E}_\sigma [\sigma_j^4] - 1) \Phi(u)^2.
\end{aligned}$$

The claim (4.1) now follows, with the choice of constant as in (4.2). $\square$

Proof of Theorem 5.6. Recall that T_J is a set of time points in $[0, T]$, indexed by an index set J with cardinality $|J| \in \mathbb{N}$. In (5.10), we observed that

$$\|G_N(u) - G(u)\| \leq |J| \sup_{0 \leq t \leq T} \|e(t; u)\|_{\ell_2^d}.$$

Fix $\rho^* > 2$. Omitting the argument u of Φ_N , Φ , G_N and G , we have

$$\begin{aligned}
\exp(\rho^* \Phi_N) &= \exp(\rho^* (\Phi_N - \Phi + \Phi)) \\
&\leq \exp(\rho^* |\Phi_N - \Phi| + \rho^* \Phi) \\
&= \exp(\rho^* |\Phi_N - \Phi|) \exp(\rho^* \Phi) \\
&\leq \exp(2\rho^* C_\Gamma (\Phi^{1/2} \|G_N - G\| + \|G - G_N\|^2)) \exp(\rho^* \Phi) \\
&\leq \frac{\exp(\rho^* \Phi)}{2} [\exp(4\rho^* C_\Gamma \Phi^{1/2} \|G_N - G\|) + \exp(4\rho^* C_\Gamma \|G - G_N\|^2)],
\end{aligned}$$

where the last two inequalities follow from (3.9) and Young's inequality $ab \leq (a^2 + b^2)/2$ for $a, b \geq 0$. Using (5.10), we therefore obtain

$$\begin{aligned}
\exp(\rho^* \Phi_N) &\leq \frac{\exp(\rho^* \Phi)}{2} \left[\exp\left(4\rho^* C_\Gamma \Phi^{1/2} |J| \sup_{0 \leq t \leq T} \|e(t)\|_{\ell_2^d}\right) \right. \\
&\quad \left. + \exp\left(4\rho^* C_\Gamma |J|^2 \sup_{0 \leq t \leq T} \|e(t)\|_{\ell_2^d}^2\right) \right],
\end{aligned}$$

where we note that we have suppressed the u -dependence of $e(t; u)$ and simply written $e(t)$. Since $\mathcal{U}$ is compact and S is continuous, it follows that G and hence Φ are continuous on $\mathcal{U}$; by the extreme value theorem, Φ is bounded on $\mathcal{U}$, i.e. $\|\Phi\|_{L_{\mu_0}^\infty(\mathcal{U})}$ is finite. Using this fact and taking expectations with respect to ν_N we obtain

$$\begin{aligned}
\mathbb{E}_{\nu_N} [\exp(\rho^* \Phi_N(u))] &\leq \frac{\exp(\rho^* \|\Phi\|_{L_{\mu_0}^\infty(\mathcal{U})})}{2} \left(\mathbb{E}_{\nu_N} \left[\exp\left(4\rho^* C_\Gamma \|\Phi\|_{L_{\mu_0}^\infty(\mathcal{U})}^{1/2} |J| \sup_{0 \leq t \leq T} \|e(t; u)\|_{\ell_2^d}\right) \right] \right. \\
&\quad \left. + \mathbb{E}_{\nu_N} \left[\exp\left(4\rho^* C_\Gamma |J|^2 \sup_{0 \leq t \leq T} \|e(t; u)\|_{\ell_2^d}^2\right) \right] \right).
\end{aligned}$$

By Corollary 5.5, the two terms on the right-hand side are finite for every $u \in \mathcal{U}$. Given the continuous dependence of the parameters of Assumptions 5.1, 5.2, and 5.3 on u , and given that $\mathcal{U}$ is a compact subset of a finite-dimensional Euclidean space, it follows that the right-hand side can be bounded by a scalar that

does not depend on any u . Hence, the function $u \mapsto \mathbb{E}_{\nu_N}[\exp(\rho^* \Phi_N(u))]$ belongs to $L_{\mu_0}^\infty(\mathcal{U}) \subset L_{\mu_0}^1(\mathcal{U})$, so that the first hypothesis of Theorem 3.9(b) holds. For the second hypothesis, observe that, since Assumption 5.3 holds for $R = +\infty$, it follows that (5.12) holds for any $n \in \mathbb{N}$, and thus (3.11) holds for any $q, s \geq 1$. Therefore the hypotheses of Theorem 3.9(b) are satisfied, and the desired conclusion follows from Theorem 3.9. $\square$