



Citation for published version:

Scheichl, R, Stuart, AM & Teckentrup, AL 2017, 'Quasi-Monte Carlo and Multilevel Monte Carlo Methods for Computing Posterior Expectations in Elliptic Inverse Problems', *SIAM/ASA Journal on Uncertainty Quantification*, vol. 5, no. 1, pp. 493-518. <https://doi.org/10.1137/16M1061692>

DOI:

[10.1137/16M1061692](https://doi.org/10.1137/16M1061692)

Publication date:

2017

Document Version

Peer reviewed version

[Link to publication](#)

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Quasi-Monte Carlo and Multilevel Monte Carlo Methods for Computing Posterior Expectations in Elliptic Inverse Problems

R. Scheichl, A.M. Stuart, A.L. Teckentrup

Abstract

We are interested in computing the expectation of a functional of a PDE solution under a Bayesian posterior distribution. Using Bayes' rule, we reduce the problem to estimating the ratio of two related prior expectations. For a model elliptic problem, we provide a full convergence and complexity analysis of the ratio estimator in the case where Monte Carlo, quasi-Monte Carlo or multilevel Monte Carlo methods are used as estimators for the two prior expectations. We show that the computational complexity of the ratio estimator to achieve a given accuracy is the same as the corresponding complexity of the individual estimators for the numerator and the denominator. We also include numerical simulations, in the context of the model elliptic problem, which demonstrate the effectiveness of the approach.

1 Introduction

Simulation frequently plays an essential role in the mathematical modelling of physical processes. However, the model parameters are often subject to uncertainty. This may be due to incomplete or inaccurate knowledge of the system or due to an inherent variability. It is important to understand how this uncertainty in the input parameters influences the reliability of the simulation outputs.

In the Bayesian framework, we initially assign a probability distribution, called the *prior distribution*, to the input parameters. In addition *observations*, related to the model outputs, are often available, and it is then possible to reduce the overall uncertainty and get a better representation of the input parameters by conditioning the prior distribution on this data. This leads to the *posterior* distribution on the input parameters. The goal of the simulations is then often to compute the expected value of a *quantity of interest* (related to the model outputs) under the posterior distribution. This is the problem of *Bayesian inference*.

Typically, the posterior distribution is intractable, in the sense that direct sampling is unavailable. One way to circumvent this problem is to use a Markov chain Monte Carlo (MCMC) approach to sample from the posterior distribution [40, 18, 11, 32, 9]. However, for large-scale applications where the number of input parameters is typically large and the solution of the forward model expensive, MCMC methods require careful tuning and may become infeasible in practice.

An alternative approach, considered in this paper, is to note that any given posterior expectation can be written as the ratio of two prior expectations, both involving the *likelihood*. The denominator is the normalising constant in Bayes' rule, namely the expected value under the prior of the likelihood. The numerator is similar, but the likelihood is weighted by the test function of interest. This approach has already been considered in [44, 41, 43], where the prior expectations are computed using adaptive sparse grid techniques. Similar ideas are also found in [1] in the context of importance sampling. The focus of this work is to use sampling methods to compute the prior

expectations, which are also well-suited to the case of high dimensional inputs. In particular, we investigate the use of Monte Carlo (MC), quasi-Monte Carlo (QMC) [38, 37, 35, 26] and multilevel Monte Carlo (MLMC) [25, 30, 2, 47, 29, 20] methods. The work is closely related to the independent works [14, 15, 23] that also investigate the use of QMC methods in computing posterior expectations. The papers [14, 15] are narrower in the range of problems considered. Their analysis is based on holomorphy arguments which require uniformly bounded coefficients, thus excluding the Gaussian case considered here. On the other hand, the range of methods considered is wider and includes higher-order QMC and multilevel QMC methods. The work [23] investigates the use of higher order QMC methods for computing posterior expectations arising from partial differential equations posed on random domains, again in the case of bounded parameters.

As a particular example, we consider the model inverse problem of determining the distribution of the diffusion coefficient of a divergence form elliptic partial differential equation (PDE) from observations of a finite set of noisy continuous functionals of the solution. The coefficient distribution is assumed to be determined by an infinite number of scalar parameters through a basis expansion. In contrast to the works [44, 41, 43, 14], our analysis includes also results in the technically demanding case of log-normal diffusion coefficients, where the differential operator depends in a non-affine way on the parameters, each of which is modelled as a Gaussian random variable under the prior distribution. We provide a full convergence and complexity analysis of the estimator of the posterior expectation in the case of MC, QMC and MLMC sampling. We also demonstrate the effectiveness of this approach for the estimation of a typical quantity of interest derived from the elliptic inverse problem. The main conclusion of our work is that, for a given accuracy, the cost of computing the posterior expectation with any of these Monte Carlo variants is proportional to the computational complexity of the same estimator for prior expectations.

The remainder of this paper is organised as follows. Section 2 provides the mathematical set-up of the inverse problem of interest, including the formulation of ratio estimators for posterior expectations. Section 3 is then devoted to the analysis of the error committed by approximating the governing equations by finite elements, and Section 4 introduces MC, QMC and MLMC estimators together with bounds on their sampling errors, extending the QMC analysis to non-linear functionals. In Section 5, we then provide a full convergence and complexity analysis of ratio estimators of posterior expectations. We demonstrate the performance of the proposed ratio estimators on a specific quantity of interest in Section 6, and finally provide some conclusions in Section 7.

2 Bayesian Inverse Problems

Let X and V be separable Banach spaces, and define the Borel measurable mappings $\mathcal{G} : X \rightarrow V$ and $\mathcal{H} : V \rightarrow \mathbb{R}^m$, for some $m \in \mathbb{N}$. We will refer to \mathcal{G} as the *forward map* and to \mathcal{H} as the *observation operator*. We denote by $\mathcal{F} : X \rightarrow \mathbb{R}^m$ the composition of \mathcal{H} and \mathcal{G} , and by $|\cdot|$ the Euclidean norm on \mathbb{R}^m . The inverse problem of interest is to determine the unknown function $u \in X$ from the noisy *observations* (or *data*) $y \in \mathbb{R}^m$ given by

$$y = \mathcal{H}(\mathcal{G}(u)) + \eta, \tag{2.1}$$

where the noise η is a realisation of the \mathbb{R}^m -valued Gaussian random variable $\mathcal{N}(0, \Gamma)$, for some (known) covariance matrix Γ . For simplicity, we will assume that $\Gamma = \sigma_\eta^2 I$, for some positive constant σ_η^2 .

We adopt a Bayesian perspective in which, in the absence of data, u is distributed according to a prior measure μ_0 . Under the conditions given in Proposition 2.1 below, the posterior distribution μ^y on the conditioned random variable $u|y$ is absolutely continuous with respect to μ_0 and given by an infinite dimensional version of Bayes' Theorem. This takes the form

$$\frac{d\mu^y}{d\mu_0}(u) = \frac{1}{Z}\theta(\mathcal{G}(u)), \quad (2.2)$$

where

$$\theta(\zeta) = \exp[-\Phi(\zeta)], \quad \Phi(\zeta) = \frac{1}{2\sigma_\eta^2} |y - \mathcal{H}(\zeta)|^2 \quad \text{and} \quad Z = \mathbb{E}_{\mu_0}[\theta(\mathcal{G}(u))]. \quad (2.3)$$

The following proposition from [45] provides conditions under which the posterior distribution μ^y is well defined and satisfies (2.2).

Proposition 2.1. *Assume the map $\mathcal{F} : X \rightarrow \mathbb{R}^m$ is continuous and $\mu_0(X) = 1$. Then the posterior distribution μ^y is absolutely continuous with respect to the prior distribution μ_0 , with Radon-Nikodym derivative given by (2.2).*

In applications, it is often of interest to compute the expectation of a functional $\phi : V \rightarrow \mathbb{R}$ of $\mathcal{G}(u)$ under the posterior distribution μ^y . If we define

$$\psi(\zeta) = \theta(\zeta) \phi(\zeta) \quad \text{and} \quad Q = \mathbb{E}_{\mu_0}[\psi(\mathcal{G}(u))], \quad (2.4)$$

it follows from (2.2) that the posterior expectation of $\phi(\mathcal{G}(u))$ can be written as

$$\mathbb{E}_{\mu^y}[\phi(\mathcal{G}(u))] = \frac{\mathbb{E}_{\mu_0}[\psi(\mathcal{G}(u))]}{\mathbb{E}_{\mu_0}[\theta(\mathcal{G}(u))]} = \frac{Q}{Z}. \quad (2.5)$$

We will approximate $\mathbb{E}_{\mu^y}[\phi(\mathcal{G}(u))]$ by using different Monte Carlo type methods to compute the prior expectations Z and Q .

2.1 Parametrisation of the Unknown Input

We consider the setting where the Banach space X is a space of real-valued functions defined on a bounded spatial domain $D \subset \mathbb{R}^d$, for some dimension $d = 1, 2$ or 3 . For ease of presentation, we shall restrict our attention to the case $X = C(\overline{D})$, the space of continuous functions on \overline{D} , but other choices are possible (cf. Remark 2.12), as is the extension to vector-valued functions.

We assume that the unknown function $u \in X$ admits a parametric representation of the form

$$u(x) = m_0(x) + \sum_{j=1}^{\infty} u_j \phi_j(x), \quad (2.6)$$

where $m_0 \in X$, $\{\phi_j\}_{j=1}^{\infty}$ denotes an infinite sequence in X (typically normalised to one in X or in a larger space containing X) and $\{u_j\}_{j=1}^{\infty} \subset \mathbb{R}^{\infty}$ denotes a set of real-valued coefficients. By randomising the coefficients $\{u_j\}_{j=1}^{\infty}$, we create real-valued random functions on \overline{D} . To this end, we introduce the deterministic, monotonically non-increasing sequence $\gamma = \{\gamma_j\}_{j=1}^{\infty}$ and the i.i.d random sequence $\xi = \{\xi_j\}_{j=1}^{\infty}$, and set $u_j = \gamma_j \xi_j$. To emphasise the dependence of u on ξ , we will write $u = u(x; \xi)$.

We will consider two specific examples of the infinite series representation (2.6), referred to as uniform priors and Gaussian priors, respectively.

2.1.1 Uniform Priors

In the case of uniform priors, we specify the i.i.d. sequence of random variables $\xi = \{\xi_j\}_{j=1}^\infty$ by choosing $\xi_j \sim U[-1, 1]$, a uniform random variable on $[-1, 1]$, and the deterministic sequence γ is chosen absolutely summable, $\gamma \in \ell^1(\mathbb{R}^\infty)$. The functions $\{\phi_j\}_{j=1}^\infty$ and m_0 are chosen as elements of $C(\overline{D})$, and are assumed normalised so that $\|\phi_j\|_{C(\overline{D})} = 1$, for all $j \in \mathbb{N}$.

We then have the following result from [12].

Lemma 2.2. *Suppose there are finite, strictly positive constants m_{\min}, m_{\max} and r such that*

$$\min_{x \in \overline{D}} m_0(x) \geq m_{\min}, \quad \max_{x \in \overline{D}} m_0(x) \leq m_{\max} \quad \text{and} \quad \|\gamma\|_{\ell^1} = \frac{r}{1+r} m_{\min}.$$

Then the following holds almost surely: the function $u(\cdot; \xi)$ defined in (2.6) is in $C(\overline{D})$ and satisfies the bounds

$$\frac{1}{1+r} m_{\min} \leq u(x; \xi) \leq m_{\max} + \frac{r}{1+r} m_{\min}, \quad \text{for almost all } x \in D.$$

Note in particular that the upper and lower bounds on u in Lemma 2.2 are independent of the particular realisation of the random sequence ξ . With $X = C(\overline{D})$, it follows from Lemma 2.2 that $\mu_0(X) = 1$. We furthermore have the following result from [12] on the spatial regularity of the function u in the case where the functions $(m_0, \{\phi_j\}_{j=1}^\infty)$ are Hölder continuous.

Lemma 2.3. *Suppose m_0 and $\{\phi_j\}_{j \geq 1}$, are in $C^\alpha(\overline{D})$, the space of Hölder continuous functions with exponent $\alpha \leq 1$, and suppose $\sum_{j=1}^\infty |\gamma_j|^2 \|\phi_j\|_{C^\alpha(\overline{D})}^\beta < \infty$, for some $\beta \in (0, 2)$. Then the function $u(\cdot; \xi)$ defined in (2.6) is in $C^t(\overline{D})$ almost surely, for any $t < \alpha\beta/2$.*

2.1.2 Gaussian Priors

For Gaussian priors, we specify the i.i.d sequence of random variables ξ by choosing $\xi_j \sim N(0, 1)$, a standard Gaussian random variable with mean 0 and variance 1. We choose the sequences $\{\phi_j\}_{j=1}^\infty$ and $\{\gamma_j^2\}_{j=1}^\infty$ to be the eigenfunctions and eigenvalues, respectively, of a covariance operator $\mathcal{C} : L^2(D) \rightarrow L^2(D)$, such that the series (2.6) is the Karhunen-Loeve (KL) expansion of the Gaussian measure $\mu_0 = N(m_0, \mathcal{C})$ on $L^2(D)$. Denote by $c : D \times D \rightarrow \mathbb{R}$ the covariance kernel corresponding to the covariance operator \mathcal{C} . It follows from Mercer's Theorem that the eigenvalues $\{\gamma_j^2\}_{j=1}^\infty$ are positive and summable, and the equality $c(x, y) = \sum_{j=1}^\infty \gamma_j^2 \phi_j(x) \phi_j(y)$ holds for almost all $x, y \in D$.

We have the following result on the spatial regularity of the function u from [4, 12].

Lemma 2.4. *Let \mathcal{C} denote the covariance operator with covariance kernel c satisfying $c(x, y) = g(\|x - y\|)$, for all $x, y \in D$, some norm $\|\cdot\|$ on \mathbb{R}^d and some Lipschitz continuous function $g \in C^{0,1}(\overline{D})$. Let $\{\phi_j\}_{j=1}^\infty$ and $\{\gamma_j^2\}_{j=1}^\infty$ be the eigenfunctions and eigenvalues of \mathcal{C} , respectively, and suppose $m_0 \in C^t(\overline{D})$, for some $t < 1/2$. Then, the function $u(\cdot; \xi)$ defined in (2.6) is also in $C^t(\overline{D})$ almost surely.*

An example of a covariance kernel $c(x, y)$ that satisfies the assumptions of Lemma 2.4 is the *exponential* covariance kernel

$$c(x, y) = \sigma^2 \exp[-\|x - y\|_r / \lambda], \tag{2.7}$$

where the positive parameters σ^2 and λ are known as the variance and correlation length, respectively, and typically $r = 1$ or 2 .

It follows from Lemma 2.4 that if the covariance operator \mathcal{C} is smooth enough, so that the function g is Lipschitz continuous, the function $u(\cdot; \xi)$ is almost surely continuous. With $X = C(\overline{D})$, it hence follows that $\mu_0(X) = 1$.

For practical applications, such as the problem described in Section 2.2, it is often of interest to construct a function that is strictly positive on \overline{D} . For this reason, we consider the function $a(\cdot; \xi) = \exp[u(\cdot; \xi)]$. Since $u(\cdot; \xi)$ is almost surely continuous, we can define almost surely the quantities

$$a_{\min}(\xi) = \min_{x \in \overline{D}} a(x; \xi), \quad \text{and} \quad a_{\max}(\xi) = \max_{x \in \overline{D}} a(x; \xi).$$

We have the following result on the boundedness of the function a [4, 5].

Lemma 2.5. *Let the assumptions of Lemma 2.4 hold. Then $a(\cdot; \xi) = \exp(u(\cdot; \xi))$ is in $C^t(\overline{D})$ almost surely, for any $t < 1/2$. Furthermore,*

$$0 < a_{\min}(\xi) \leq a(x; \xi) \leq a_{\max}(\xi) < \infty, \quad \text{for almost all } x \in D \text{ and } \xi \in \mathbb{R}^\infty,$$

and $a_{\min}^{-1} \in L^r(\mathbb{R}^\infty)$, $a_{\max} \in L^r(\mathbb{R}^\infty)$ and $a \in L^r(\mathbb{R}^\infty, C^t(\overline{D}))$, for all $r \in [1, \infty)$.

Other, smoother covariance kernels, such as the *Gaussian* kernel

$$c(x, y) = \sigma^2 \exp[-\|x - y\|_2^2 / \lambda^2]$$

or the kernels from the Matérn family, also satisfy the assumptions of Lemmas 2.4 and 2.5, but they lead to a significantly higher spatial regularity $t \geq 1/2$ of a .

2.1.3 Finite-dimensional approximation

In simulations, it is often necessary to use a finite-dimensional approximation of the unknown u . Given the parametrisation (2.6), this can be achieved by simply truncating the series at finite truncation order J , or through best N -term approximations [7, 48]. For simplicity, we here choose the former, and make the following assumption in the remainder of this paper.

Assumption A1. (*Finite Truncation Order*) Suppose the coefficients $\{\gamma_j\}_{j=J+1}^\infty$ are all equal to zero, for some finite $J \in \mathbb{N}$, such that

$$u(x; \xi) = u(x; \xi_J) = m_0(x) + \sum_{j=1}^J \gamma_j \xi_j \phi_j(x), \quad (2.8)$$

where $\xi_J := \{\xi_j\}_{j=1}^J \in \mathbb{R}^J$.

An important question is how one should optimally choose the truncation order J in (2.8), and the answer typically involves a trade-off between choosing J sufficiently large to retain a required accuracy, and sufficiently small to avoid an unnecessarily large computational cost associated to sampling from u . We will in this paper assume that $J \in \mathbb{N}$ is given, and will not explicitly discuss how to choose J . We refer the interested reader to the works [4, 35, 47, 26].

The series (2.8) defines a linear mapping $P : \mathbb{R}^J \rightarrow X$, with $P(\xi_J) = u$, and we will define the prior measure μ_0 on X as the pushforward under P of a suitable measure \mathbb{P} defined on the

coefficient space \mathbb{R}^J , equipped with the Borel product σ -algebra. For the uniform priors considered in section 2.1.1, the measure \mathbb{P} is the product measure

$$\mathbb{P}(d\xi_{\underline{J}}) = \prod_{j=1}^J \frac{d\xi_j}{2}. \quad (2.9)$$

For the Gaussian priors in section 2.1.2, we have

$$\mathbb{P}(d\xi_{\underline{J}}) = \prod_{j=1}^J \frac{1}{\sqrt{2\pi}} \exp[-\xi_j^2/2] d\xi_j. \quad (2.10)$$

Corollary 2.6. *Suppose Assumption A1 holds. Then (i) Lemmas 2.2 and 2.3 hold; and (ii) if $\{\phi_j\}_{j=1}^J$ and $\{\gamma_j^2\}_{j=1}^J$ are chosen as the first J elements of $\{\phi_j\}_{j=1}^\infty$ and $\{\gamma_j^2\}_{j=1}^\infty$ in Lemma 2.4, respectively, then Lemma 2.4 and 2.5 hold.*

Proof. Part (i) follows directly from Lemmas 2.2 and 2.3, since the parametrisation (2.8) is just a special case of (2.6). Part (ii) is proven in [4]. \square

Remark 2.7. (Alternative Approximations) The truncated parametrisation (2.8) is not the only way to obtain an approximation of u from which we can easily produce samples for simulation. In the case of Gaussian priors, knowledge of the covariance kernel c allows us to assemble the covariance matrix of the Gaussian vector $[u(x_1), u(x_2), \dots, u(x_n)]$, for any $n \in \mathbb{N}$ and $\{x_i\}_{i=1}^n \subseteq D$, and we can hence use methods based on factorisations of the covariance matrix, such as [17], to sample from u at a finite number of locations in the domain D . In applications such as the elliptic problem discussed in section 2.2, this is usually sufficient, since typically quadrature methods are used to compute the numerical approximation discussed in section 3. For more details, we refer the interested reader to [27, 28, 46].

2.2 Model Elliptic Problem

We consider the model inverse problem of determining the distribution of the diffusion coefficient of a divergence form elliptic partial differential equation (PDE) from observations of a finite set of noisy continuous functionals of the solution. Let $D \subset \mathbb{R}^d$, for $d = 1, 2$ or 3 , be a bounded Lipschitz domain, and denote by ∂D its boundary. The forward problem which underlies the inverse problem of interest here is to find the solution $p(\cdot; \xi_{\underline{J}})$ of the following linear elliptic PDE,

$$-\nabla \cdot (k(x; \xi_{\underline{J}}) \nabla p(x; \xi_{\underline{J}})) = f(x) \quad \text{in } D, \quad p(\cdot; \xi_{\underline{J}}) = 0 \quad \text{on } \partial D, \quad (2.11)$$

for given functions $k(\cdot; \xi_{\underline{J}}) \in C(\overline{D})$ and $f \in H^{-1}(D)$. Although all results in this section apply also in the case of infinite-dimensional parameter vectors ξ , we restrict our attention to finite-dimensional $\xi_{\underline{J}}$ for consistency.

The variational formulation of (2.11) is to find $p(\cdot; \xi_{\underline{J}}) \in H_0^1(D)$ such that

$$b(p, q; \xi_{\underline{J}}) = L(q), \quad \text{for all } q \in H_0^1(D), \quad (2.12)$$

where the bilinear form b and the linear functional L are defined as usual, for all $v, w \in H_0^1(D)$ by

$$b(v, w; \xi_{\underline{J}}) = \int_D k(x; \xi_{\underline{J}}) \nabla v(x) \cdot \nabla w(x) dx \quad \text{and} \quad L(w) = \langle f, w \rangle_{H^{-1}(D), H_0^1(D)}. \quad (2.13)$$

We say that $p(\cdot; \xi_{\underline{J}})$ is a weak solution to (2.11) iff $p(\cdot; \xi_{\underline{J}}) \in H_0^1(D)$ and $p(\cdot; \xi_{\underline{J}})$ satisfies (2.12).

In the inverse problem, we take the coefficient k to be a function of the unknown u , in which case both the coefficient k and the solution p depend on the random sequence $\xi_{\underline{J}}$. When the dependence on $\xi_{\underline{J}}$ of k and p is irrelevant, we will simply write $k = k(x)$ and $p = p(x)$. With the unknown function $u(\cdot; \xi_{\underline{J}})$ as in Section 2.1, we choose

- $k(\cdot; \xi_{\underline{J}}) = u(\cdot; \xi_{\underline{J}})$, in the case of the uniform priors described in Section 2.1.1, and
- $k(\cdot; \xi_{\underline{J}}) = k^* + \exp(u(\cdot; \xi_{\underline{J}}))$, in the case of Gaussian priors, for some given continuous non-negative function $k^* \geq 0$.

By Lemmas 2.2 and 2.5, both these choices ensure that the diffusion coefficient $k(\cdot; \xi_{\underline{J}})$ in (2.12) is strictly positive on D , \mathbb{P} -almost surely. In terms of the notation in previous sections, the Banach space X is the space of continuous functions $C(\overline{D})$ as before. The forward map \mathcal{G} is defined by $\mathcal{G}(u) = p$, i.e. it maps the unknown function $u(\cdot; \xi_{\underline{J}})$ to the solution $p(\cdot; \xi_{\underline{J}})$. (Note that the definition of \mathcal{G} differs between the two choices $k = u$ and $k = k^* + \exp(u)$.) We take the Banach space V as the Sobolev space $H_0^1(D)$.

Existence and uniqueness of the weak solution $p(\cdot; \xi_{\underline{J}})$ is ensured by the Lax-Milgram Theorem. As in previous sections, let $k_{\min}(\xi_{\underline{J}})$ and $k_{\max}(\xi_{\underline{J}})$ be such that

$$0 < k_{\min}(\xi_{\underline{J}}) \leq k(x; \xi_{\underline{J}}) \leq k_{\max}(\xi_{\underline{J}}) < \infty, \quad \text{for almost all } x \in D \text{ and for } \xi_{\underline{J}} \text{ } \mathbb{P}\text{-almost surely.}$$

For uniform priors, $k_{\min}(\xi_{\underline{J}})$ and $k_{\max}(\xi_{\underline{J}})$ are independent of $\xi_{\underline{J}}$. If $k^*(x) > 0$, for all $x \in \overline{D}$, then $k_{\min}(\xi_{\underline{J}})$ is also independent of $\xi_{\underline{J}}$ in the Gaussian case.

Definition 2.8. We will refer to the coefficient k as *uniformly elliptic* (respectively *uniformly bounded*) when $k_{\min}(\xi_{\underline{J}})$ (respectively $k_{\max}(\xi_{\underline{J}})$) is independent of $\xi_{\underline{J}}$.

The following is a direct consequence of the Lax-Milgram Lemma and Lemmas 2.2 and 2.5.

Lemma 2.9. For \mathbb{P} -almost all $\xi_{\underline{J}} \in \mathbb{R}^J$, there exists a unique weak solution $p(\cdot; \xi_{\underline{J}}) \in H_0^1(D)$ to the variational problem (2.12) and

$$|p(\cdot; \xi_{\underline{J}})|_{H^1(D)} \leq \frac{\|f\|_{H^{-1}(D)}}{k_{\min}(\xi_{\underline{J}})}.$$

Furthermore, $p \in L_{\mathbb{P}}^r(\mathbb{R}^J, H_0^1(D))$, for all $r \in [1, \infty)$. If k is uniformly elliptic, then the result holds also for $r = \infty$.

In order to conclude on the well-posedness of the posterior distribution μ^y , we furthermore have the following result on the continuity of the forward map \mathcal{G} .

Lemma 2.10. The map $\mathcal{G} : X \rightarrow V$, $\mathcal{G}(u) = p$, is continuous.

Proof. Denote by p_1 and p_2 two weak solutions of (2.12) with the same right hand side f and with coefficients k_1 and k_2 , respectively. Let k_{\min} and k_{\max} be such that

$$0 < k_{\min} \leq k_i(x) \leq k_{\max} < \infty, \quad \text{for almost all } x \in D,$$

for $i = 1, 2$. Then it follows from the variational formulation (2.12) that

$$|p_1 - p_2|_{H^1(D)} \leq \frac{\|f\|_{H^{-1}(D)}}{k_{\min}^2} \|k_1 - k_2\|_{C(\overline{D})}.$$

In the case $k = u$, the continuity of \mathcal{G} now follows immediately. In the case $k = \exp(u)$, the continuity of \mathcal{G} follows from the continuity of the exponential function. \square

We then have the following corollary to Proposition 2.1, which follows immediately from Lemmas 2.2, 2.4 and 2.10, together with the continuity of the observation operator \mathcal{H} .

Corollary 2.11. *For the forward map \mathcal{G} defined by $\mathcal{G}(u) = p$, the posterior measure μ^y is absolutely continuous with respect to the prior measure μ_0 , with Radon-Nikodym derivative (2.2).*

Remark 2.12. (Piecewise continuous coefficients) Although we here restrict our attention to the case of continuous random coefficients, the theory extends to the piecewise continuous case where a further source of randomness can be introduced in the partitioning of the computational domain D into sub-domains. The well-posedness of the posterior distribution in this case was shown in [?]. The regularity and spatial discretisation error (as discussed in Section 3) were analysed in [47, 46].

3 Finite Element Discretisation

In this section, we analyse the error introduced in the computation of the prior expectations Z and Q by a finite element approximation of the forward map \mathcal{G} . We consider only standard, continuous, piecewise linear finite elements on polygonal/polyhedral domains in detail. To this end, denote by $\{\mathcal{T}_h\}_{h>0}$ a shape-regular family of simplicial triangulations of the Lipschitz polygonal/polyhedral domain D , parametrised by their mesh width $h := \max_{\tau \in \mathcal{T}_h} \text{diam}(\tau)$. Associated with each triangulation \mathcal{T}_h we define the space

$$V_h := \{q_h \in C(\overline{D}) : q_h|_{\tau} \text{ linear for all } \tau \in \mathcal{T}_h \text{ and } q_h|_{\partial D} = 0\} \quad (3.1)$$

of continuous, piecewise linear functions on D that vanish on the boundary ∂D .

The finite element approximation to (2.12), denoted by p_h , is now the unique function in V_h that satisfies

$$b(p_h, q_h; \xi_{\mathcal{J}}) = L(q_h), \quad \text{for all } q_h \in V_h, \quad (3.2)$$

where the bilinear form b and the functional L are as in (2.13). Note that, in particular, this implies that p_h satisfies the same bound as in Lemma 2.9:

$$|p_h(\cdot; \xi_{\mathcal{J}})|_{H^1(D)} \leq \|f\|_{H^{-1}(D)} / k_{\min}(\xi_{\mathcal{J}}). \quad (3.3)$$

The approximate forward map $\mathcal{G}_h : X \rightarrow V$ is then defined by $\mathcal{G}_h(u) = p_h$, and we denote the resulting approximations of Z and Q , respectively, by

$$Z_h = \mathbb{E}_{\mu_0}[\theta(\mathcal{G}_h(u))] \quad \text{and} \quad Q_h = \mathbb{E}_{\mu_0}[\psi(\mathcal{G}_h(u))].$$

A standard technique to prove convergence of finite element approximations of functionals is to use a duality argument, similar to the classic Aubin Nitsche trick used to prove optimal convergence rates for the L^2 norm. In the context of the elliptic PDE (2.11) with random coefficients, this analysis was performed in [47]. We here summarise the main results of the error analysis, and show that if the observation operator \mathcal{H} and the functional of interest ϕ are smooth enough, the finite element error in the prior expectations Z_h and Q_h converges at the optimal rate.

Let $v, w \in H_0^1(D)$. Given a functional $F : H_0^1(D) \rightarrow \mathbb{R}$, we denote by $D_v F(w)$ its Fréchet derivative at w , applied to v . With p and p_h as before, we define

$$\begin{aligned} \overline{D_v F}(p, p_h) &= \int_0^1 D_v F(p + \lambda(p_h - p)) \, d\lambda, \\ \text{and } |\overline{D_v F}|(p, p_h) &= \int_0^1 |D_v F(p + \lambda(p_h - p))| \, d\lambda, \end{aligned}$$

which in some sense are averaged derivatives of F on the path from p to p_h . Let us now define the following dual problem: find $z \in H_0^1(D)$ such that

$$b(q, z; \xi_{\underline{J}}) = \overline{D_q F}(p, p_h), \quad \text{for all } q \in H_0^1(D). \quad (3.4)$$

Denote the finite element approximation of the dual solution z by $z_h \in V_h$. It then follows from the Fundamental Theorem of Calculus, Galerkin orthogonality of the primal problem (2.12) and boundedness of the bilinear form b that

$$|F(p) - F(p_h)| = |b(p - p_h, z; \xi_{\underline{J}})| = |b(p - p_h, z - z_h; \xi_{\underline{J}})| \leq k_{\max}(\xi_{\underline{J}}) |p - p_h|_{H^1(D)} |z - z_h|_{H^1(D)}.$$

In order to prove convergence of the finite element error $|F(p) - F(p_h)|$, it hence suffices to prove convergence of $|p - p_h|_{H^1(D)}$ and $|z - z_h|_{H^1(D)}$. For our further analysis, we make the following assumption on the smoothness of the maps ϕ and \mathcal{H} . Examples of functionals satisfying Assumption A2 are discussed in [47], and include linear functionals, powers of linear functionals and boundary fluxes.

Assumption A2. (*Differentiability*) Let ϕ and \mathcal{H}_i , $i = 1, \dots, m$, be continuously Fréchet differentiable on the path $\{p + \lambda(p - p_h)\}_{\lambda \in [0,1]}$, and suppose that there exist $t_* \in [0, 1]$, $q_* \in [1, \infty]$ and $C_\phi, C_{\mathcal{H}} \in L_{\mathbb{P}}^{q_*}(\mathbb{R}^J)$ such that $f \in H^{t_*-1}(D)$,

$$\overline{D_v \phi}(p, p_h) \leq C_\phi(\xi_{\underline{J}}) \|v\|_{H^{1-t_*}(D)}, \quad \text{and} \quad \overline{D_v \mathcal{H}_i}(p, p_h) \leq C_{\mathcal{H}}(\xi_{\underline{J}}) \|v\|_{H^{1-t_*}(D)},$$

for all $v \in H_0^1(D)$ and almost all $\xi_{\underline{J}} \in \mathbb{R}^J$.

Let now $F = \phi$ or $F = \mathcal{H}_i$, for some $i \in \{1, \dots, m\}$. To get well-posedness of the primal problem (2.12) and the dual problem (3.4), as well as existence and uniqueness of the solutions $p(\cdot; \xi_{\underline{J}}) \in H_0^1(D)$ and $z(\cdot; \xi_{\underline{J}}) \in H_0^1(D)$, for almost all $\xi_{\underline{J}} \in \mathbb{R}^J$, it is sufficient to assume that Assumption A2 holds with $t_* = 0$. However, in order to prove convergence of the finite element approximations, it is necessary to require stronger spatial regularity of p and z , which requires Assumption A2 to hold for some $t_* > 0$. We have the following result from [5, 47]. The assumptions on D being polygonal and convex are purely to simplify the presentation. Proposition 3.1 also holds for piecewise smooth or for non-convex domains, but typically with stronger restrictions on the range of s .

Proposition 3.1. *Let D be a Lipschitz polygonal, convex domain, let $k \in L_{\mathbb{P}}^{r_*}(\mathbb{R}^J, C^t(\overline{D}))$ and $k_{\min}^{-1} \in L_{\mathbb{P}}^{r_*}(\mathbb{R}^J)$, for some $t \in (0, 1]$ and $r_* \in [1, \infty]$, and let Assumption A2 hold with $t_* = t$ and $q_* = r_*$. Then, the solutions p and z of (2.12) and (3.4) are both in $L_{\mathbb{P}}^r(\mathbb{R}^J, H^{1+s}(D))$, for any $s < t$ and $r < r_*$. The result also holds for $r = r_*$ or for $s = t$, if $r_* = \infty$ or $t = 1$, respectively.*

We will now show that the functionals θ and ψ appearing in the prior expectations Z and Q , respectively, satisfy the bounds in Assumption A2 provided ϕ and \mathcal{H} satisfy Assumption A2, as well as the growth conditions in Assumption A3 below.

Assumption A3. (*Boundedness*) Suppose there are constants $M_1, M_2 > 0$ and $n_1, n_2 \in \mathbb{N}$, such that

$$|\phi(v)| \leq M_1 \left(1 + |v|_{H^1(D)}^{n_1}\right) \quad \text{and} \quad |\mathcal{H}(v)| \leq M_2 \left(1 + |v|_{H^1(D)}^{n_2}\right), \quad \text{for all } v \in H_0^1(D). \quad (3.5)$$

Recall the product and chain rules for Fréchet derivatives for functionals $F_1, F_2 : H_0^1(D) \rightarrow \mathbb{R}$ and a function $f : \mathbb{R} \rightarrow \mathbb{R}$:

$$D_v(F_1 F_2)(w) = F_2(w) D_v F_1(w) + F_1(w) D_v F_2(w) \quad \text{and} \quad D_v(f \circ F_1)(w) = D_{D_v F_1(w)} f(F_1(w)).$$

We then have the following result.

Lemma 3.2. *Let Assumption A2 hold with $t_* \in [0, 1]$ and $q_* \in [1, \infty]$, and suppose Assumption A3 holds. Then*

$$\overline{|D_v \theta|}(p, p_h) \leq C_\theta(\xi_{\underline{J}}) \|v\|_{H^{1-t_*}(D)}, \quad \text{and} \quad \overline{|D_v \psi|}(p, p_h) \leq C_\psi(\xi_{\underline{J}}) \|v\|_{H^{1-t_*}(D)},$$

for all $v \in H_0^1(D)$ and for $\xi_{\underline{J}} \mathbb{P}$ -almost surely, where $C_\theta(\xi_{\underline{J}})$ and $C_\psi(\xi_{\underline{J}})$ are in $L_{\mathbb{P}}^r(\mathbb{R}^J)$, for all $r \in [1, q_*)$. If k is uniformly elliptic, then the result holds also for $r = q_*$.

Proof. First, we use the chain rule and product rule for Fréchet derivatives to obtain

$$D_v \theta(w) = D_v \left(\exp \left[-\frac{1}{2\sigma_\eta^2} \sum_{i=1}^m (y_i - \mathcal{H}_i(w))^2 \right] \right) = -\theta(w) \frac{1}{\sigma_\eta^2} \sum_{i=1}^m D_v \mathcal{H}_i(w) (y_i - \mathcal{H}_i(w)).$$

Denoting $p_\lambda = p + \lambda(p_h - p)$, we then have

$$\overline{|D_v \theta|}(p, p_h) = \int_0^1 |D_v \theta(p_\lambda)| \, d\lambda \leq \int_0^1 \theta(p_\lambda) \frac{1}{\sigma_\eta^2} \sum_{i=1}^m |D_v \mathcal{H}_i(p_\lambda)| |y_i - \mathcal{H}_i(p_\lambda)| \, d\lambda.$$

By the definition of θ in (2.3), we have $\theta(p_\lambda) \leq 1$, for all $\lambda \in [0, 1]$. By Assumption A3, it follows that

$$|y_i - \mathcal{H}_i(p_\lambda)| \leq C(1 + |p_\lambda|_{H^1(D)}^{n_2}) \leq C(1 + \|f\|_{H^{-1}(D)}^{n_2} k_{\min}^{-n_2}(\xi_{\underline{J}})),$$

for some (generic) constant C independent of the mesh size h and random parameter $\xi_{\underline{J}}$. It then follows that

$$\overline{|D_v \theta|}(p, p_h) \leq \frac{C}{\sigma_\eta^2} (1 + \|f\|_{H^{-1}(D)}^{n_2} k_{\min}^{-n_2}(\xi_{\underline{J}})) \sum_{i=1}^m \overline{|D_v \mathcal{H}_i|}(p, p_h).$$

With $C_\theta(\xi_{\underline{J}}) = \frac{mC}{\sigma_\eta^2} (1 + \|f\|_{H^{-1}(D)}^{n_2} k_{\min}^{-n_2}(\xi_{\underline{J}})) C_{\mathcal{H}}(\xi_{\underline{J}})$, it then follows from Assumption A2 that

$$\overline{|D_v \theta|}(p, p_h) \leq C_\theta(\xi_{\underline{J}}) \|v\|_{H^{1-t_*}(D)}.$$

Next, using the product rule for Fréchet derivatives, together with the result just proved, we have

$$D_v \psi(w) = \theta(w) D_v \phi(w) - \phi(w) \theta(w) \frac{1}{\sigma_\eta^2} \sum_{i=1}^m D_v \mathcal{H}_i(w) (\delta_i - \mathcal{H}_i(w)).$$

With p_λ as before, it then follows that

$$\begin{aligned} \overline{|D_v \psi|}(p, p_h) &= \int_0^1 |D_v \psi(p_\lambda)| \, d\lambda \\ &\leq \int_0^1 \theta(p_\lambda) \left[|D_v \phi(p_\lambda)| + \phi(p_\lambda) \frac{1}{\sigma_\eta^2} \sum_{i=1}^m |D_v \mathcal{H}_i(p_\lambda)| |\delta_i - \mathcal{H}_i(p_\lambda)| \right] \, d\lambda. \end{aligned}$$

Now $\theta(p_\lambda) \leq 1$, for all $\lambda \in [0, 1]$. By Assumption A3, it follows that

$$|\phi(p_\lambda)| \max_{i \in \{1, \dots, m\}} |y_i - \mathcal{H}_i(p_\lambda)| \leq C(1 + |p_\lambda|_{H^1(D)}^{n_2})^2 \leq C(1 + \|f\|_{H^{-1}(D)}^{2n_2} k_{\min}^{-2n_2}(\xi_{\underline{J}})),$$

for some (generic) constant C independent of the mesh size h and random parameter $\xi_{\underline{J}}$. It then follows that

$$\overline{|D_v \psi|}(p, p_h) \leq \overline{|D_v \phi|}(p, p_h) + \frac{C}{\sigma_\eta^2} (1 + \|f\|_{H^{-1}(D)}^{2n_2} k_{\min}^{-2n_2}(\xi_{\underline{J}})) \sum_{i=1}^m \overline{|D_v \mathcal{H}_i|}(p, p_h).$$

With $C_\psi(\xi_{\underline{J}}) = C_\phi(\xi_{\underline{J}}) + \frac{mC}{\sigma_\eta^2} (1 + \|f\|_{H^{-1}(D)}^{2n_2} k_{\min}^{-2n_2}(\xi_{\underline{J}})) C_{\mathcal{H}}(\xi_{\underline{J}})$, it then follows that

$$\overline{|D_v \psi|}(p, p_h) \leq C_\psi(\xi_{\underline{J}}) \|v\|_{H^{1-t_*}(D)}.$$

Finally, recall that by Lemmas 2.2 and 2.5, we have $k_{\min}^{-1}(\xi_{\underline{J}}) \in L_{\mathbb{P}}^\infty(\mathbb{R}^J)$ if k is uniformly elliptic, and $k_{\min}^{-1}(\xi_{\underline{J}}) \in L_{\mathbb{P}}^q(\mathbb{R}^J)$, for any $1 \leq q < \infty$, otherwise. Hence, it follows from Assumption A2, together with Hölder's and Minkowski's inequalities, that $C_\theta(\xi_{\underline{J}})$ and $C_\psi(\xi_{\underline{J}})$ are in $L_{\mathbb{P}}^r(\mathbb{R}^J)$, for all $r \in [1, q_*)$. If k is uniformly elliptic, we can also set $r = q_*$. \square

Bounds on the finite element errors $|\theta(p) - \theta(p_h)|$ and $|\psi(p) - \psi(p_h)|$ now follow directly from Proposition 3.1 and Lemma 3.2.

Theorem 3.3. *Under the assumptions of Proposition 3.1 and Lemma 3.2 with $t \in (0, 1]$ and $r_* \in [1, \infty]$, we have*

$$\|\theta(p) - \theta(p_h)\|_{L_{\mathbb{P}}^r(\mathbb{R}^J)} \leq C_{k,f,\theta,D} h^{2s}, \quad \text{and} \quad \|\psi(p) - \psi(p_h)\|_{L_{\mathbb{P}}^r(\mathbb{R}^J)} \leq C_{k,f,\psi,D} h^{2s},$$

for any $s < t$ and $r < r_*$. The constants $C_{k,f,\theta,D}$ and $C_{k,f,\psi,D}$ are independent of h . If $r_* = \infty$, we can also bound the $L_{\mathbb{P}}^\infty$ norms. If $t = 1$, we can set $s = 1$.

Proposition 3.1, Lemma 3.2 and Theorem 3.3 hold, without any additional assumptions, also for infinite-dimensional parameter vectors $\xi \in \mathbb{R}^\infty$ [5, 47].

4 Sampling methods

In this section, we briefly recall the main ideas behind Monte Carlo (MC), Multilevel Monte Carlo (MLMC) and quasi-Monte Carlo (QMC) estimators to compute the prior expectation $Q_h = \mathbb{E}_{\mu_0}[\psi(p_h)]$. The estimators for $Z_h = \mathbb{E}_{\mu_0}[\theta(p_h)]$ are defined analogously. We also provide bounds on the sampling error of the estimators, which will become useful for bounding the mean square error in Section 5. For more details, we refer the reader to [40, 6, 35, 26].

4.1 Monte Carlo estimators

The standard Monte Carlo estimator for Q_h is

$$\widehat{Q}_{h,N}^{\text{MC}} = \frac{1}{N} \sum_{i=1}^N \psi(p_h(\cdot; \xi_{\underline{J}}^{(i)})), \quad (4.1)$$

where $\xi_{\underline{J}}^{(i)}$ is the i th sample of $\xi_{\underline{J}}$ from the distribution \mathbb{P} , and N independent samples are computed in total. The estimator (4.1) is an unbiased estimator of Q_h , with variance

$$\mathbb{V}[\widehat{Q}_{h,N}^{\text{MC}}] = \frac{\mathbb{V}[\psi(p_h)]}{N}. \quad (4.2)$$

4.2 Multilevel Monte Carlo estimators

The main idea of multilevel Monte Carlo estimation is simple. Linearity of the expectation operator implies that

$$\mathbb{E}_{\mu_0}[\psi(p_h)] = \mathbb{E}_{\mu_0}[\psi(p_{h_0})] + \sum_{\ell=1}^L \mathbb{E}_{\mu_0}[\psi(p_{h_\ell}) - \psi(p_{h_{\ell-1}})],$$

where $\{h_\ell\}_{\ell=0}^L$ are the mesh widths of a sequence of increasingly fine triangulations \mathcal{T}_{h_ℓ} with $h_L = h$, the finest mesh width, and $k_1 \leq h_{\ell-1}/h_\ell \leq k_2$, for all $\ell = 1, \dots, L$ and some $1 < k_1 \leq k_2 < \infty$. The multilevel idea is now to estimate each of the terms independently using a Monte Carlo estimator. Setting for convenience $Y_0^\psi = \psi(p_{h_0})$, and $Y_\ell^\psi = \psi(p_{h_\ell}) - \psi(p_{h_{\ell-1}})$, for $\ell = 1, \dots, L$, we define the MLMC estimator as

$$\widehat{Q}_{h,\{N_\ell\}}^{\text{ML}} = \sum_{\ell=0}^L \widehat{Y}_{\ell,N_\ell}^{\psi,\text{MC}} = \sum_{\ell=0}^L \frac{1}{N_\ell} \sum_{i=1}^{N_\ell} Y_\ell^\psi(\cdot; \xi_{\underline{J}}^{(i,\ell)}), \quad (4.3)$$

where importantly the quantity $Y_\ell^\psi(\cdot; \xi_{\underline{J}}^{(i,\ell)})$ uses the same sample $\xi_{\underline{J}}^{(i,\ell)}$ on both meshes. The estimator (4.3) is an unbiased estimator of Q_h , with variance

$$\mathbb{V}[\widehat{Q}_{h,\{N_\ell\}}^{\text{ML}}] = \sum_{\ell=0}^L \frac{\mathbb{V}[Y_\ell^\psi]}{N_\ell} \leq C \sum_{\ell=0}^L \frac{h_\ell^{4s}}{N_\ell}, \quad (4.4)$$

where the last inequality follows from Theorem 3.3, with a constant C independent of $\{h_\ell\}_{\ell=0}^L$ and with $0 \leq s < t \leq 1$, as defined in Proposition 3.1.

When defining the MLMC estimator (4.3), one can in fact also use level-dependent truncation levels J_ℓ . This approach was analysed in [47], and can lead to further significant gains in terms of computational cost.

4.3 Quasi-Monte Carlo estimators

Quasi-Monte Carlo methods are classically formulated as quadrature rules over the unit cube $[0, 1]^J$, for some $J \in \mathbb{N}$. Treating $\xi_{\underline{J}}$ as a deterministic parameter vector distributed according to the product uniform or Gaussian measure, respectively,

$$\mathbb{E}_{\mu_0}[\psi(p_h)] = \int_{[0,1]^J} \psi(p_h(\cdot; (\Phi_J^{-1}(v)))) dv, \quad (4.5)$$

where $\xi_{\underline{J}} = \Phi_J^{-1}(v)$ denotes the inverse cumulative normal applied to each entry of v in the Gaussian case. In the uniform case, Φ_J^{-1} is the simple change of variables mapping v_j to $2v_j - 1$. We will use a randomly shifted lattice rule to approximate the integral (4.5). This takes the form

$$\widehat{Q}_{h,N}^{\text{QMC}} = \frac{1}{N} \sum_{i=1}^N \psi\left(p_h(\cdot; \tilde{\xi}_{\underline{J}}^{(i)})\right), \quad \text{where } \tilde{\xi}_{\underline{J}}^{(i)} := \Phi_J^{-1}\left(\text{frac}\left(\frac{iz}{N} + \Delta\right)\right), \quad (4.6)$$

$z \in \{1, \dots, N-1\}^J$ is a *generating vector*, Δ is a uniformly distributed *random shift* on $[0, 1]^J$, and "frac" denotes the fractional part function, applied component-wise. To ensure that every one-dimensional projection of the lattice rule has N distinct values we furthermore assume that each component z_j of z satisfies $\text{gcd}(z_j, N) = 1$ (cf. [16]).

The variance of the QMC estimator (4.6) is given by

$$\mathbb{V}[\widehat{Q}_{h,N}^{\text{QMC}}] = \mathbb{E}_{\Delta}[(\mathbb{E}_{\mu_0}[\psi(p_h)] - \widehat{Q}_{h,N}^{\text{QMC}})^2]. \quad (4.7)$$

To bound it, we make the following assumption on the integrand $\psi(p_h)$.

Assumption A4. Let $c_1 > 0$ be a constant independent of J and let $b_j := \gamma_j \|\phi_j\|_{C^0(\overline{D})}$, for $j \in \mathbb{N}$. We assume that, for any multi-index $\nu \in \{0, 1\}^J$ with $|\nu| = \sum_{j \leq J} \nu_j$,

$$\left| \frac{\partial^{|\nu|} \psi(p_h)}{\partial \xi_{\underline{J}}^{\nu}} \right| \leq C_{k,f,\psi,D} \frac{c_1^{|\nu|} |\nu|!}{k_{\min}(\xi_{\underline{J}})} \prod_{j=1}^J b_j^{\nu_j}.$$

For linear functionals ψ on $H_0^1(D)$, this has been proved in [35] and in [26] for the uniform and the Gaussian cases, respectively. In both cases, we can choose $c_1 = 1/\ln 2$. However, in the Bayesian setting, both ψ and θ are inherently non-linear functionals of p . Nevertheless, if ϕ and \mathcal{H} are linear functionals of p , and k is uniformly elliptic, Assumption A4 can be proved by using the classical Faà di Bruno formula [10], a multidimensional version of the chain rule. A proof for θ in the case $m = 1$ can be found in Appendix A. We omit the proof for ψ or for $m > 1$. A proof for general analytic functionals ϕ and \mathcal{H} of p would be even more technical and require the use of generalisations of Faà di Bruno's formula to Fréchet derivatives. For this reason, we simply work under Assumption A4.

In the case of uniform priors, Assumption A4 was proven to hold in [14] using arguments from complex analysis and the holomorphy of p_h as a function of $\xi_{\underline{J}}$.

Lemma 4.1. *Suppose Assumption A4 holds and the sequence $\{b_j\}_{j=1}^{\infty}$ is in $l^q(\mathbb{R}^{\infty})$, for some $q \in (0, 1]$. Then, a randomly shifted lattice rule can be constructed via a component-by-component algorithm in $\mathcal{O}(JN \log N)$ cost, such that*

$$\mathbb{V}[\widehat{Q}_{h,N}^{\text{QMC}}] \leq \begin{cases} C_{\psi,q,\delta} N^{-1/\delta}, & \text{if } q \in (0, 2/3], \\ C_{\psi,q} N^{-(1/q-1/2)}, & \text{if } q \in (2/3, 1), \end{cases}$$

for any $\delta \in (1/2, 1]$, independently of J . For $q = 1$ and under further assumptions given in [26, Theorem 20] and [35, Theorem 6.4], we have $\mathbb{V}[\widehat{Q}_{h,N}^{\text{QMC}}] \leq C_{\psi} N^{-1/2}$.

Proof. The proof follows those in [35, 26] with suitable changes to the product and order dependent (POD) weights if $c_1 \neq 1/\ln 2$. \square

It is even possible to combine quasi-Monte Carlo sampling and multilevel estimation and the gains are complementary [36, 34], but we will not include these estimators or their analysis here.

5 Mean Square Error and Computational Complexity

We will now use the results from Sections 3 and 4 to bound the mean square error (MSE) of estimators for the ratio Q/Z . To this end, let us denote by \widehat{Z}_h and \widehat{Q}_h one of the Monte Carlo type estimators discussed in Section 4 for Z_h and Q_h , respectively. Let us define the mean square error

$$e\left(\widehat{Q}_h/\widehat{Z}_h\right)^2 = \mathbb{E}\left[\left(\frac{Q}{Z} - \frac{\widehat{Q}_h}{\widehat{Z}_h}\right)^2\right]. \quad (5.1)$$

For MC and MLMC estimators, the expectation in the expression (5.1) above is with respect to the prior measure μ_0 on X . For QMC estimators, it is with respect to the random shift Δ .

Rearranging the mean square error and using the triangle inequality, we have

$$\begin{aligned} e\left(\widehat{Q}_h/\widehat{Z}_h\right)^2 &= \frac{1}{Z^2} \mathbb{E}\left[\left(Q - \widehat{Q}_h + (\widehat{Q}_h/\widehat{Z}_h)(\widehat{Z}_h - Z)\right)^2\right] \\ &\leq \frac{2}{Z^2} \left(\mathbb{E}\left[(\widehat{Q}_h - Q)^2\right] + \mathbb{E}\left[(\widehat{Q}_h/\widehat{Z}_h)^2(\widehat{Z}_h - Z)^2\right]\right). \end{aligned} \quad (5.2)$$

Our further analysis depends on the integrability of $\widehat{Q}_h/\widehat{Z}_h$, and we thus consider separately the cases of uniformly and non-uniformly elliptic coefficients k .

5.1 Uniformly elliptic case

In the case where the coefficient k is uniformly elliptic, we have the following result on the integrability of $\widehat{Q}_h/\widehat{Z}_h$. The assumptions on ϕ and \mathcal{H} are more general than Assumption A3, and allow for very general non-linear growth. Since, in general, $\widehat{Z}_{h,\{N_\ell\}}^{\text{ML}}$ could be negative, we require stronger assumptions in the case of multilevel Monte Carlo estimators. In particular, we require the assumptions of Theorem 3.3 to hold with $r = \infty$, which means that the coefficient k needs to be uniformly bounded as well as uniformly elliptic. The analysis of MLMC in Lemma 5.1 below therefore does not apply in the case of Gaussian priors.

Lemma 5.1. *Suppose k is uniformly elliptic and there are two constants $M_1, M_2 > 0$, such that*

$$|\phi(v)| \leq M_1 \quad \text{and} \quad |\mathcal{H}(v)| \leq M_2, \quad \text{for all } v \in H_0^1(D) \quad \text{with } |v|_{H^1(D)} \leq \|f\|_{H^{-1}(D)}/k_{\min}.$$

Then $\widehat{Q}_{h,N}^{\text{MC}}/\widehat{Z}_{h,N}^{\text{MC}} \in L_{\mathbb{P}}^\infty(\mathbb{R}^J)$, and $\widehat{Q}_{h,N}^{\text{QMC}}/\widehat{Z}_{h,N}^{\text{QMC}} \in L_{\Delta}^\infty([0,1]^J)$, with L^∞ -norms bounded independently of h and N .

If in addition h_0 is sufficiently small and the assumptions of Theorem 3.3 hold with $r = \infty$, we also have $\widehat{Q}_{h,\{N_\ell\}}^{\text{ML}}/\widehat{Z}_{h,\{N_\ell\}}^{\text{ML}} \in L_{\mathbb{P}}^\infty(\mathbb{R}^J)$, with L^∞ -norm bounded independently of $\{h_\ell\}$, $\{N_\ell\}$ and L .

Proof. Using the definition of $\widehat{Q}_{h,N}^{\text{MC}}$ in (4.1), as well as the bound in (3.3) and the fact that $\theta(v) \leq 1$, for all $v \in H_0^1(D)$, it follows that

$$|\widehat{Q}_{h,N}^{\text{MC}}| = \left| \frac{1}{N} \sum_{i=1}^N \phi(p_h(\cdot; \xi_{\underline{J}}^{(i)})) \theta(p_h(\cdot; \xi_{\underline{J}}^{(i)})) \right| \leq M_1$$

and

$$\widehat{Z}_{h,N}^{\text{MC}} = \frac{1}{N} \sum_{i=1}^N \theta(p_h(\cdot; \xi_{\underline{J}}^{(i)})) \geq \exp\left(-\frac{|y|^2 + mM_2^2}{2\sigma_\eta^2}\right) =: b > 0.$$

Since the upper bound on $\widehat{Q}_{h,N}^{\text{MC}}$ and the lower bound on $\widehat{Z}_{h,N}^{\text{MC}}$ are independent of the random samples $\{\xi_{\underline{J}}^{(i)}\}_{i=1}^N$, the claim of the lemma follows for Monte Carlo estimators. The proof for quasi-Monte Carlo estimators is identical.

For multilevel Monte Carlo estimators, an upper bound on $\widehat{Q}_{h,\{N_\ell\}}^{\text{ML}}$ follows as before. On the other hand, to bound $\widehat{Z}_{h,\{N_\ell\}}^{\text{ML}}$ we can use Theorem 3.3 which implies that

$$\widehat{Z}_{h,\{N_\ell\}}^{\text{ML}} \geq b - \sum_{\ell=1}^L Ch_\ell^{2s},$$

for a constant C independent of $\{h_\ell\}$. If we choose h_0 sufficiently small such that $\sum_{\ell=1}^\infty h_\ell^{2s} < b/C$, this lower bound on $\widehat{Z}_{h,\{N_\ell\}}^{\text{ML}}$ is positive and independent of $\{h_\ell\}$, $\{N_\ell\}$ and L . The claim of the Lemma then follows also for MLMC estimators. \square

Using Lemma 5.1 and Hölder's inequality, it then follows from (5.2) that

$$e\left(\widehat{Q}_h/\widehat{Z}_h\right)^2 \leq 2/Z^2 \max\{1, \|\widehat{Q}_h/\widehat{Z}_h\|_{L^\infty}^2\} \left(\mathbb{E}\left[(Q - \widehat{Q}_h)^2\right] + \mathbb{E}\left[(Z - \widehat{Z}_h)^2\right]\right).$$

Thus, the MSE of the ratio $\widehat{Q}_h/\widehat{Z}_h$ can be bounded by the sum of the MSEs of \widehat{Q}_h and \widehat{Z}_h . Using the fact that, for Q_h the mean of the estimator \widehat{Q}_h ,

$$\mathbb{E}\left[(Q - \widehat{Q}_h)^2\right] = (\mathbb{E}[Q - Q_h])^2 + \mathbb{V}\left[\widehat{Q}_h\right] \quad (5.3)$$

and the results from Sections 3 and 4, this gives the following bounds on the MSEs.

Theorem 5.2. *Suppose the relevant assumptions of Proposition 3.1, Lemma 4.1 and Lemma 5.1 hold in each case. Then*

$$\begin{aligned} e\left(\widehat{Q}_{h,N}^{\text{MC}}/\widehat{Z}_{h,N}^{\text{MC}}\right)^2 &\leq C_{\text{MC}}(N^{-1} + h^{4s}), \\ e\left(\widehat{Q}_{h,N}^{\text{QMC}}/\widehat{Z}_{h,N}^{\text{QMC}}\right)^2 &\leq C_{\text{QMC}}(N^{-1/\delta} + h^{4s}), \\ e\left(\widehat{Q}_{h,\{N_\ell\}}^{\text{ML}}/\widehat{Z}_{h,\{N_\ell\}}^{\text{ML}}\right)^2 &\leq C_{\text{ML}}\left(\sum_{\ell=0}^L \frac{h_\ell^{4s}}{N_\ell} + h^{4s}\right), \end{aligned}$$

for some $1/2 < \delta \leq 1$ and for some $0 < s \leq 1$, related to the spatial regularity of the data (cf. Proposition 3.1), and for constants C_{MC} , C_{QMC} and C_{ML} independent of h , N , $\{h_\ell\}$, $\{N_\ell\}$ and L .

We note that the convergence rates of the mean square errors in Theorem 5.2 are identical to the convergence rates obtained for the individual prior estimators \widehat{Q}_h and \widehat{Z}_h .

5.2 Non-uniformly elliptic case

If the functionals ϕ and \mathcal{H} are uniformly bounded, in the sense that

$$|\phi(v)| \leq M_1 \quad \text{and} \quad |\mathcal{H}(v)| \leq M_2, \quad \text{for all } v \in H_0^1(D), \quad \text{and} \quad |\theta(p) - \theta(p_h)| \leq M_3 h, \quad (5.4)$$

for some constants $M_1, M_2, M_3 > 0$, then the analysis in Lemma 5.1 carries over to the non-uniformly elliptic case, with only minor modifications in the proof.

For more general functionals ϕ and \mathcal{H} , the analysis is significantly more difficult and we are only able to analyse Monte Carlo estimators. We restrict to functionals ϕ and \mathcal{H} that satisfy the polynomial growth conditions in Assumption A4. Then we can follow an approach similar to [19] to obtain the following integrability result on $\widehat{Q}_{h,N}^{\text{MC}}/\widehat{Z}_{h,N}^{\text{MC}}$.

Lemma 5.3. *Suppose that ϕ and \mathcal{H} satisfy Assumption A4 and that the same N i.i.d. samples $\{\xi_J^{(i)}\}_{i=1}^N$ are used in the estimators $\widehat{Q}_{h,N}^{\text{MC}}$ and $\widehat{Z}_{h,N}^{\text{MC}}$. Then $\widehat{Q}_{h,N}^{\text{MC}}/\widehat{Z}_{h,N}^{\text{MC}} \in L_{\mathbb{P}}^r(\mathbb{R}^J)$ and $\|\widehat{Q}_{h,N}^{\text{MC}}/\widehat{Z}_{h,N}^{\text{MC}}\|_{L_{\mathbb{P}}^r(\mathbb{R}^J)}$ can be bounded independent of h and N , for all $1 \leq r < \infty$.*

Proof. To simplify the presentation, let us denote $p_h^i = p_h(\cdot; \xi_J^{(i)})$. Then, due to (3.3) and (3.5) we can bound $\theta(p_h^i) > 0$, \mathbb{P} -almost surely. Hence, it follows from (4.1) that

$$\left| \frac{\widehat{Q}_{h,N}^{\text{MC}}}{\widehat{Z}_{h,N}^{\text{MC}}} \right| = \left| \frac{\frac{1}{N} \sum_{i=1}^N \phi(p_h^i) \theta(p_h^i)}{\frac{1}{N} \sum_{i=1}^N \theta(p_h^i)} \right| \leq \max_{1 \leq i \leq N} |\phi(p_h^i)|. \quad (5.5)$$

Using the same argument as in the proof of [19, Lemma 1], for any convex and non-decreasing function $\rho : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, we have

$$\rho \left(\mathbb{E} \left[\max_{1 \leq i \leq N} |\phi(p_h^i)|^r \right] \right) \leq \mathbb{E} \left[\rho \left(\max_{1 \leq i \leq N} |\phi(p_h^i)|^r \right) \right] \leq \sum_{i=1}^N \mathbb{E} \left[\rho \left(|\phi(p_h^i)|^r \right) \right].$$

Choosing $\rho(x) = |x|^{\tilde{r}/r}$, for some $r \leq \tilde{r} < \infty$, and using (3.5) we have

$$\left\| \max_{1 \leq i \leq N} |\phi(p_h^i)| \right\|_{L_{\mathbb{P}}^{\tilde{r}}(\mathbb{R}^J)} \leq \|\phi(p_h)\|_{L_{\mathbb{P}}^{\tilde{r}}(\mathbb{R}^J)} N^{1/\tilde{r}} \leq M_1 \left(1 + \|f\|_{H^{-1}(D)}^{n_2} \|k_{\min}^{-1}\|_{L_{\mathbb{P}}^{n_2 \tilde{r}}(\mathbb{R}^J)}^{n_2} \right) N^{1/\tilde{r}},$$

The term in the bracket is finite due to Lemma 2.5 and the claim of the Lemma now follows if we choose $\tilde{r} \geq \ln N$. \square

Theorem 5.4. *Suppose the assumptions of Proposition 3.1 and Lemma 5.3 hold. Then*

$$e \left(\widehat{Q}_{h,N}^{\text{MC}} / \widehat{Z}_{h,N}^{\text{MC}} \right)^2 \leq C_{\text{MC}} (N^{-1} + h^{4s}),$$

for some $0 < s \leq 1$ related to the spatial regularity of the data (cf. Proposition 3.1) and for a constant $C_{\text{MC}} > 0$ independent of h and N .

Proof. Since $\widehat{Q}_{h,N}^{\text{MC}}/\widehat{Z}_{h,N}^{\text{MC}}$ is not in $L_{\mathbb{P}}^{\infty}(\mathbb{R}^J)$ in this case, we apply the Cauchy-Schwarz inequality to the second term on the right hand side of (5.2) to obtain

$$e \left(\widehat{Q}_{h,N}^{\text{MC}} / \widehat{Z}_{h,N}^{\text{MC}} \right)^2 \leq \frac{2}{Z^2} \left(\mathbb{E} \left[(\widehat{Q}_{h,N}^{\text{MC}} - Q)^2 \right] + \|\widehat{Q}_{h,N}^{\text{MC}} / \widehat{Z}_{h,N}^{\text{MC}}\|_{L_{\mathbb{P}}^4(\mathbb{R}^J)}^2 \|\widehat{Z}_{h,N}^{\text{MC}} - Z\|_{L_{\mathbb{P}}^4(\mathbb{R}^J)}^2 \right). \quad (5.6)$$

To bound $\|\widehat{Z}_{h,N}^{\text{MC}} - Z\|_{L_{\mathbb{P}}^4(\mathbb{R}^J)}$ we apply the triangle inequality and consider separately $\|Z_h - Z\|_{L_{\mathbb{P}}^4(\mathbb{R}^J)}$ and $\|\widehat{Z}_{h,N}^{\text{MC}} - Z_h\|_{L_{\mathbb{P}}^4(\mathbb{R}^J)}$. The former is bounded by Ch^{2s} due to Theorem 3.3.

To bound the latter, let $X_i := \theta(p_h(\cdot; \xi_J^{(i)}))$, $i = 1, \dots, N$. Since the range of θ is $[0, 1]$, this is a sequence of i.i.d. random variables with finite mean $m := Z_h$, finite variance $\sigma_X^2 := \mathbb{E}[(\theta(p_h) - Z_h)^2]$ and finite central fourth moment $\tau_X^4 := \mathbb{E}[(\theta(p_h) - Z_h)^4]$. A direct calculation gives

$$\|\widehat{Z}_{h,N}^{\text{MC}} - Z_h\|_{L_{\mathbb{P}}^4(\mathbb{R}^J)}^4 = \mathbb{E} \left[\left(\widehat{X}_N - m \right)^4 \right] = \frac{3(N-1)\sigma_X^4/2 + \tau_X^4}{N^3} \leq \frac{3\sigma_X^4/2 + \tau_X^4/N}{N^2}.$$

Thus, the result follows from (5.6) together with Theorem 3.3 and Lemma 5.3. \square

The proofs of Lemma 5.3 and Theorem 5.4 can potentially also be extended to the case of Quasi-Monte Carlo estimators. However, it seems impossible to satisfy Assumption A4 in the case of non-uniformly elliptic coefficients k and so we did not pursue this any further.

The analysis in the case of multilevel Monte Carlo estimators is complicated by the fact that the multilevel estimator $\widehat{Z}_{h,\{N_\ell\}}^{\text{ML}}$ can take on negative values. This precludes the approach in the proof of Lemma 5.3. The existence of moments of a ratio of random variables where the denominator is not strictly positive has been the subject of research since the 1930s [3, 21, 31, 39, 33, 22], and is a problem not yet fully solved. A possible approach to show existence of moments of the ratio $\widehat{Q}_{h,\{N_\ell\}}^{\text{ML}}/\widehat{Z}_{h,\{N_\ell\}}^{\text{ML}}$ could be to use the Central Limit Theorem in [8], which shows that the individual MLMC estimators are asymptotically normally distributed as the number of levels and the number of samples per level tend to infinity. Hinkley [31] then gives an explicit expression of the cumulative distribution function of the ratio of two correlated normal random variables, together with its limiting normal distribution, as the denominator tends to a normal random variable with non-zero mean and zero variance.

5.3 Computational ε -cost

Based on the bounds on the mean square errors given in Theorems 5.2 and 5.4, we now analyse the computational complexity of the various estimators of our quantity of interest Q/Z . We are interested in bounding the ε -cost, i.e., the cost required to achieve a MSE of order ε^2 . Since the convergence rates of the mean square error are the same as for the individual estimators \widehat{Q}_h and \widehat{Z}_h , bounds on the computational ε -cost can be proved as in [47, 26].

We denote by C_ℓ the cost of obtaining one sample of $\theta(p_{h_\ell})$ and/or $\psi(p_{h_\ell})$. This cost will typically also depend on the truncation parameter J , but we will not make this dependence explicit here. We furthermore denote by \mathcal{C}_{MC} , \mathcal{C}_{ML} and \mathcal{C}_{QMC} the computational cost of the ratio estimator $\widehat{Q}_h/\widehat{Z}_h$ based on MC, MLMC and QMC estimators, respectively.

Theorem 5.5. *Let the conclusions of Theorem 5.2 or Theorem 5.4 hold and suppose*

$$C_\ell \leq C_\gamma h_\ell^{-\gamma}, \quad \text{for some } \gamma > 0.$$

Then for any $\varepsilon < e^{-1}$, there exist a constant $C^{\text{ML}} > 0$, a value $L \in \mathbb{N}$ and a sequence $\{N_\ell\}_{\ell=0}^L$, such that

$$e \left(\widehat{Q}_{h,\{N_\ell\}}^{\text{ML}} / \widehat{Z}_{h,\{N_\ell\}}^{\text{ML}} \right)^2 \leq \varepsilon^2 \quad \text{and} \quad \mathcal{C}_{\text{ML}} \leq \begin{cases} C^{\text{ML}} \varepsilon^{-2}, & \text{if } s < \gamma/4, \\ C^{\text{ML}} \varepsilon^{-2} (\log \varepsilon)^2, & \text{if } s = \gamma/4, \\ C^{\text{ML}} \varepsilon^{-\gamma/2s}, & \text{if } s > \gamma/4, \end{cases}$$

where $0 < s \leq 1$ is related to the spatial regularity of the data (cf. Proposition 3.1). Furthermore, there exist positive constants C^{MC} and C^{QMC} and values of h and N , such that

$$e\left(\widehat{Q}_{h,N}^{\text{MC}}/\widehat{Z}_{h,N}^{\text{MC}}\right)^2 \leq \varepsilon^2 \quad \text{and} \quad \mathcal{C}_{\text{MC}} \leq C^{\text{MC}} \varepsilon^{-2-\gamma/2s},$$

$$e\left(\widehat{Q}_{h,N}^{\text{QMC}}/\widehat{Z}_{h,N}^{\text{QMC}}\right)^2 \leq \varepsilon^2 \quad \text{and} \quad \mathcal{C}_{\text{QMC}} \leq C^{\text{QMC}} \varepsilon^{-2\delta-\gamma/2s}, \quad \text{for some } 1/2 < \delta \leq 1.$$

Theorem 5.5 shows that MLMC and QMC can outperform standard MC in terms of the growth rate of the ε -cost. However, both QMC and MLMC also require stronger assumptions than MC (cf section 4).

6 Numerical examples

We now study the performance of the ratio estimators on a typical model problem. As the forward model, we take the elliptic equation

$$-\nabla \cdot (k(x; \xi_{\underline{J}})) \nabla p(x; \xi_{\underline{J}}) = 0, \quad \text{in } D = (0, 1)^2, \quad (6.1)$$

subject to the deterministic, mixed boundary conditions $p|_{x_1=0} = 1, p|_{x_1=1} = 0$ and zero Neumann conditions on the remainder of the boundary. The prior distribution on the coefficients $\xi_{\underline{J}}$ is Gaussian, as in Section 2.1.2, and k is a (truncated) log-normal random field. We choose the exponential covariance function (2.7) with $r = 1$, correlation length $\lambda = 0.3$ and variance $\sigma^2 = 1$. The mean m_0 is chosen to be 0. In this case, the assumptions of Proposition 3.1 hold for any $t < 1/2$.

For the spatial discretisation, we use standard, continuous, piecewise linear finite elements on a uniform triangular mesh. The stiffness matrix is assembled using the trapezoidal rule for quadrature. The mesh hierarchy for the MLMC estimator is generated by uniform refinement of a uniform grid with coarsest mesh width $h_0 = 1/8$, and $h_{\ell-1}/h_\ell = 2$, for all $\ell = 1, \dots, L$.

The quantity of interest ϕ is the outflow over the boundary at $x_1 = 1$. To obtain optimal convergence rates of the finite element error, we compute $\phi(p_h)$ as

$$\phi(p_h) = - \int_D k(x; \xi_{\underline{J}}) \nabla w_h(x) \cdot \nabla p_h(x; \xi_{\underline{J}}) dx,$$

for a suitably chosen weight function w_h with $w_h|_{x_1=0} = 0, w_h|_{x_1=1} = 1$ [47]. In particular, we choose $w_h \in V_h$ to be one at the nodes of the finite element mesh on the boundary $x_1 = 1$ and zero at all other nodes.

The data y is generated from the solution of equation (6.1) with a random sample $\xi_{\underline{J}}$ from the prior distribution, on a fine reference mesh with $h^* = 1/256$. The observation functional \mathcal{H} is taken as a local average pressure, representing a regularised point evaluation. To obtain m -dimensional data y , we take the uniform finite element mesh on $[0, 1]^2$ with grid size $1/(\sqrt{m} + 1)$, and evaluate the local average pressure at the m interior nodes in this mesh. The average is taken over the six elements of the finite element mesh with $h^* = 1/256$ adjacent to that node. We furthermore add observational noise to the data y , which is a realisation of an m -dimensional normal random variable with mean zero and covariance $\sigma_\eta^2 I$.

To generate samples of k , we use a truncated Karhunen-Loève expansion [24], i.e. we truncate the infinite expansion (2.6) at a finite order $J_{KL} = 1400$. An alternative that allows to sample from the infinite expansion (2.6) would be the circulant embedding method [17, 27, 28].

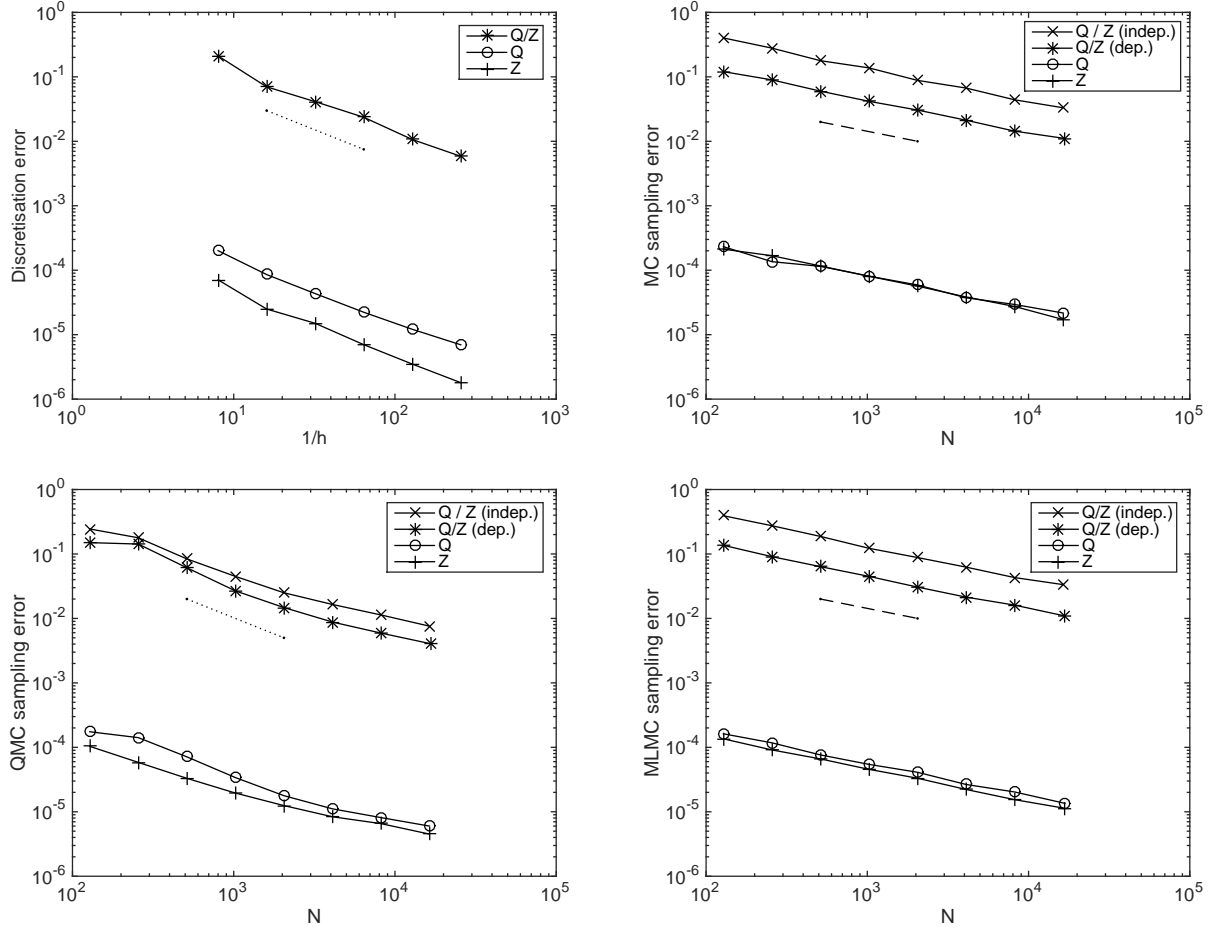


Figure 1: Convergence w.r.t. h of the discretisation errors $|Q_h/Z_h - Q_{2h}/Z_{2h}|$, $|Q_h - Q_{2h}|$ and $|Z_h - Z_{2h}|$ (top left), as well as convergence w.r.t. N of the sampling errors $\mathbb{E}[(\hat{Q}_h/\hat{Z}_h - Q_h/Z_h)^2]^{1/2}$, $\mathbb{E}[(\hat{Q}_h - Q_h)^2]^{1/2}$ and $\mathbb{E}[(\hat{Z}_h - Z_h)^2]^{1/2}$ for MC (top right), QMC (bottom left) and MLMC (bottom right), respectively. The dotted and dashed reference slopes are -1 and $-1/2$, respectively.

For the QMC estimators, we choose a lattice rule with product weight parameters $\gamma_j = 1/j^2$ and one random shift. The generating vector for the rule used is available from Frances Kuo's website (<http://web.maths.unsw.edu.au/~fkuo/>) as "lattice-39102-1024-1048576.3600". We point out here that this generating vector is a standard, off the shelf generating vector, rather than a generating vector specifically constructed for the weights implicitly defined in Assumption A4 and Lemma 4.1. In practice we found this generating vector to work well, even though the convergence rates in Lemma 4.1 were not proven for this particular choice.

6.1 Mean Square Error

We start by investigating the discretisation error, the sampling error and the mean square error of the ratio estimators, for a fixed number of observations $m = 9$ and a fixed level of observational noise $\sigma_\eta^2 = 0.09$.




Figure 2: Computational cost of ratio estimators $\widehat{Q}_h/\widehat{Z}_h$ to achieve a RMSE $e(\widehat{Q}_h/\widehat{Z}_h)$ of ε , using the same random samples in \widehat{Q}_h and in \widehat{Z}_h (left) and using different random samples in \widehat{Q}_h and in \widehat{Z}_h (right), respectively.

Figure 1 shows the discretisation error and the sampling errors of the different estimators. The top left plot shows the discretisation error $|Q_h/Z_h - Q_{2h}/Z_{2h}|$, as well as the individual discretisation errors $|Q_h - Q_{2h}|$ and $|Z_h - Z_{2h}|$. We see these errors decay linearly in h , as predicted by Theorem 3.3.

The other three plots show the sampling error $\mathbb{E}[(\widehat{Q}_h/\widehat{Z}_h - Q_h/Z_h)^2]^{1/2}$, as well as the individual sampling errors $\mathbb{E}[(\widehat{Q}_h - Q_h)^2]^{1/2}$ and $\mathbb{E}[(\widehat{Z}_h - Z_h)^2]^{1/2}$, for MC (top right), QMC (bottom left) and MLMC (bottom right). The mesh size h is fixed at $h = 1/16$, and the “exact” expected values Q_h and Z_h are estimated with MLMC with a very large number of samples. For MC and QMC, N on the horizontal axis represents the number of samples. For MLMC, N represents the equivalent number of solves on the finest grid $h = 1/16$ that would lead to the same cost as the MLMC estimator. This means that for a given N , the cost of all three estimators is the same. The number of samples N_ℓ in the MLMC estimator was chosen proportional to $h_\ell^{-(4s+\gamma)/2} \approx h_\ell^{-2}$, as suggested by the optimisation in [25, 6], assuming $s \approx 1/2$ and $\gamma \approx 2$. We show results for ratio estimators with the same random samples used in \widehat{Q}_h and \widehat{Z}_h , referred to as dependent estimators, as well as ratio estimators with different random samples used in \widehat{Q}_h and \widehat{Z}_h , referred to as independent estimators. For MC and MLMC, we observe a convergence rate of $N^{-1/2}$. For QMC, we observe a convergence rate which is significantly faster than order $N^{-1/2}$ and almost order N^{-1} .

Figure 2 compares the computational costs of the different estimators to achieve an RMSE of ε . The computational cost of the estimators was computed as Nh^{-2} for the MC and QMC estimators, and as $N_0h_0^{-2} + \sum_{\ell=1}^L N_\ell(h_\ell^{-2} + h_{\ell-1}^{-2})$ for the MLMC estimator. The bias $|Q_h/Z_h - Q/Z|$ was estimated from the values of $|Q_h/Z_h - Q_{2h}/Z_{2h}|$ shown in Figure 1. As predicted by Theorem 5.2 with $\gamma \approx 2$ and $s \approx 1/2$, the cost of the MC estimator grows with about ε^{-4} , the cost of the QMC estimator grows with about ε^{-3} , and the cost of the MLMC estimator grows with about ε^{-2} .

6.2 Dependency on m and σ_η^2

Finally, we look at the dependency of the sampling error on the number of observations m and on the noise level σ_η^2 . For large values of m and small values of σ_η^2 , we expect the posterior distribution




Figure 3: Sampling errors $\mathbb{E}[(\widehat{Q}_h/\widehat{Z}_h - Q_h/Z_h)^2]^{1/2}$ as a function of noise level σ_η^2 (left) and as a function of number of observations m (right), respectively.

μ^y to concentrate on a small region of the parameter space. The ratio estimators sample from the prior distribution and do not make use of this fact. We expect the sampling errors to grow with increasing m and decreasing σ_η^2 . To ameliorate this problem, one can under certain assumptions rescale the parameter space before applying the ratio estimators, see e.g. [42] for details.

Figure 3 shows the sampling error of ratio estimators based on using the same samples in \widehat{Q}_h and \widehat{Z}_h . We observe a mild growth of the sampling errors both with increasing m and decreasing σ_η^2 , but the growth is not dramatic and all estimators appear to be robust over a large range of practically interesting values. The fact that the sampling error for MLMC based estimators grows more quickly than for MC and QMC based estimators, is at least partly caused by the fact that $\mathbb{V}[\psi_{h_0}]$ and $\mathbb{V}[\psi_{h_1} - \psi_{h_0}]$ become of the same size for small σ_η^2 or large m , making the choice $N_\ell = Ch_\ell^{-2}$ less and less optimal. Experiments with estimators based on using independent samples in \widehat{Q}_h and \widehat{Z}_h also showed growth of sampling errors for small σ_η^2 and large m , in fact much faster than in the case of dependent estimators.

7 Conclusions and further work

In Bayesian inverse problems, the goal is often to compute the expected value of a quantity of interest under the posterior distribution. For sampling based approaches, one has to overcome the difficulty that the posterior distribution is typically intractable, in the sense that direct sampling from it is unavailable since the normalisation constant is unknown. We considered here an approach based on Bayes' theorem that computes an estimate of the normalisation constant and estimates the posterior expectation as the ratio of two prior expectations. To compute the prior expectations, we considered the sampling based approaches of Monte Carlo, quasi-Monte Carlo and multilevel Monte Carlo estimators. For a model elliptic inverse problem, we provided a full convergence and complexity analysis of the resulting ratio estimators. Our theory shows that asymptotically the complexity of computing the posterior expectation with this approach is the same as computing prior expectations, and this result is also confirmed numerically for a typical model problem in

uncertainty quantification.

It would be interesting to compare the performance of the ratio estimators considered in this work to Markov chain Monte Carlo (MCMC) and multilevel Markov chain Monte Carlo (MLMCMC) methods [32, 18]. Especially in the case of small noise level σ_η^2 or large number of observations m , MCMC based approaches might explore the posterior distribution more efficiently. In terms of the ε -cost of estimators, the analysis and simulations in [18] show that the computational cost of a standard MCMC estimator grows at the same rate as a ratio estimator based on MC, and the cost of an MLMCMC estimator will grow at the same rate as a ratio estimator based on MLMC. The constants appearing in these estimates, for MCMC based approaches, depend on quantities like the acceptance rate, autocorrelation and possibly the dimension of the parameter space. For the high dimensional problems considered in this work, these constants might be very large.

A Proof of Assumption A4 for θ for linear and scalar \mathcal{H}

Let $m = 1$, let \mathcal{H} be a linear functional on $V = H_0^1(D)$ and let $k_{\min}^* := \min_{x \in \overline{D}} k^*(x) > 0$ and thus $k_{\min}(\xi_{\underline{J}}) \geq k_{\min}^* > 0$, for all $\xi_{\underline{J}} \in \mathbb{R}^J$. Let us assume without loss of generality that $y = 0$ and $\sigma_\eta^2 = 1$, and for simplicity let $\xi = \xi_{\underline{J}} \in \mathbb{R}^J$. Then, $\theta(p_h(\cdot; \xi)) = g(h(\xi))$ with $g(\zeta) := \exp(-\zeta^2/2) \leq 1$ and $h(\xi) := \mathcal{H}(p_h(\cdot; \xi))$. To simplify the presentation, we write $g_n := \frac{d^n g}{d\zeta^n}(h(\xi))$ and $h_\mu := \frac{\partial^{|\mu|} h}{\partial \xi^\mu}(\xi)$ where μ is a multi-index in $\{0, 1\}^J$.

Let C be a generic constant independent of J , ν and ξ . First note that due to (3.3) and the linearity of \mathcal{H}

$$|h(\xi)| \leq \|\mathcal{H}\|_{H^{-1}(D)} \|p_h\|_{H_0^1(D)} \leq \frac{\|\mathcal{H}\|_{H^{-1}(D)} \|f\|_{H^{-1}(D)}}{k_{\min}^*} =: \kappa_*.$$

Then, we have $g_n(\zeta) = (-1)^n H_n(\zeta) g(\zeta)$ where H_n is the n th Hermite polynomial, and so

$$|g_n(h(\xi))| \leq C \max\{1, |h(\xi)|^n\} g(h(\xi)) \leq C \max\{1, \kappa_*^n\}. \quad (\text{A.1})$$

Moreover, it was shown in [26, Theorem 16] that, for linear \mathcal{H} ,

$$|h_\mu(\xi)| \leq \kappa_* \frac{|\mu|!}{(\ln 2)^{|\mu|}} \prod_{j=1}^J b_j^{\mu_j}. \quad (\text{A.2})$$

Now, Faa di Bruno's formula for the special case where $\nu \in \{0, 1\}^J$ and where $h(\xi)$ is scalar (cf. [10, Corollary 2.10]) states that

$$\theta_\nu = \sum_{r=1}^{|\nu|} g_r \sum_{P(r, \nu)} \prod_{i=1}^r h_{\mu^{(i)}}$$

where

$$P(r, \nu) := \left\{ \mu^{(1)}, \dots, \mu^{(r)} : \mathbf{0} \prec \mu^{(1)} \prec \dots \prec \mu^{(r)} \text{ and } \sum_{i=1}^r \mu^{(i)} = \nu \right\}$$

and \prec indicates some unique linear ordering of multi-indices (see [10, p. 505] for an example). And so, using (A.1) and (A.2), we get

$$\begin{aligned} |\theta_{\boldsymbol{\nu}}| &\leq C \left| \sum_{r=1}^{|\boldsymbol{\nu}|} \max\{1, \kappa_*^r\} \sum_{P(r, \boldsymbol{\nu})} \prod_{i=1}^r \left(\kappa_* \frac{|\boldsymbol{\mu}^{(i)}|!}{(\ln 2)^{|\boldsymbol{\mu}^{(i)}|}} \prod_{j=1}^J b_j^{\mu_j^{(i)}} \right) \right| \\ &\leq C \left(\frac{\max\{\kappa_*, \kappa_*^2\}}{\ln 2} \right)^{|\boldsymbol{\nu}|} \underbrace{\left(\sum_{r=1}^{|\boldsymbol{\nu}|} \sum_{P(r, \boldsymbol{\nu})} \prod_{i=1}^r |\boldsymbol{\mu}^{(i)}|! \right)}_{=:\rho_{\boldsymbol{\nu}}} \prod_{j=1}^J b_j^{\nu_j} \end{aligned}$$

since all elements of $P(r, \boldsymbol{\nu})$ satisfy $\sum_{i=1}^r \boldsymbol{\mu}^{(i)} = \boldsymbol{\nu}$ and $\sum_{i=1}^r |\boldsymbol{\mu}^{(i)}| = |\boldsymbol{\nu}|$. It remains to bound $\rho_{\boldsymbol{\nu}}$. We give a simple but fairly crude bound.

First note that, for each element $(\boldsymbol{\mu}^{(1)}, \dots, \boldsymbol{\mu}^{(r)}) \in P(r, \boldsymbol{\nu})$, the moduli $|\boldsymbol{\mu}^{(i)}|$, $i = 1, \dots, r$, form a partition of $n := |\boldsymbol{\nu}|$. Hence, instead of partitioning the summands in $\rho_{\boldsymbol{\nu}}$ into the subsets $P(r, \boldsymbol{\nu})$, we can also sum over all possible partitions k_1, \dots, k_r of n with $1 \leq r \leq n$. The partition function $p(n)$ is the number of possible partitions of n . It is bounded by $\exp(\pi\sqrt{2n/3})$ [13]. For each partition k_1, \dots, k_r of n , the number of possible elements in $P(r, \boldsymbol{\nu})$ that satisfy $|\boldsymbol{\mu}_i| = k_i$ can be bounded by

$$\binom{n}{k_1} \binom{n-k_1}{k_2} \binom{n-\sum_{i=1}^2 k_i}{k_3} \dots \binom{n-\sum_{i=1}^{r-1} k_i}{k_r} = \frac{n!}{k_1! k_2! \dots k_r!}$$

Elements of $P(r, \boldsymbol{\nu})$ where $|\boldsymbol{\mu}_i| = |\boldsymbol{\mu}_j|$, for some $i \neq j$, are counted twice in this bound. Since $\prod_{i=1}^r |\boldsymbol{\mu}^{(i)}|! = k_1! k_2! \dots k_r!$, we finally get the bound

$$\rho_{\boldsymbol{\nu}} \leq p(|\boldsymbol{\nu}|) |\boldsymbol{\nu}|! \leq \exp\left(\pi\sqrt{\frac{2|\boldsymbol{\nu}|}{3}}\right) |\boldsymbol{\nu}|!.$$

Hence, there exists a constant $c_p > 1$ such that Assumption A4 holds with $c_1 := c_p \max\{\kappa_*, \kappa_*^2\} / \ln 2$.

References

- [1] S. AGAPIOU, O. PAPASPILIOPOULOS, D. SANZ-ALONSO, AND A. M. STUART, *Importance sampling: computational complexity and intrinsic dimension*. Available as arXiv preprint arXiv:1511.06196.
- [2] A. BARTH, CH. SCHWAB, AND N. ZOLLINGER, *Multilevel Monte Carlo finite element method for elliptic PDE's with stochastic coefficients*, Numer. Math., 119 (2011), pp. 123–161.
- [3] R.C. CEARY, *The frequency distribution of the quotient of two normal variables*, J. Roy. Statist. Soc., 93 (1930), pp. 442–446.
- [4] J. CHARRIER, *Strong and weak error estimates for the solutions of elliptic partial differential equations with random coefficients*, SIAM J. Numer. Anal., 50 (2012), pp. 216–246.
- [5] J. CHARRIER, R. SCHEICHL, AND A. L. TECKENTRUP, *Finite element error analysis of elliptic PDEs with random coefficients and its application to multilevel Monte Carlo methods*, SIAM J. Numer. Anal., 51 (2013), pp. 322–352.
- [6] K. A. CLIFFE, M. B. GILES, R. SCHEICHL, AND A. L. TECKENTRUP, *Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients*, Comput. Vis. Sci., 14 (2011), pp. 3–15.

- [7] A. COHEN, R. DEVORE, AND CH. SCHWAB, *Convergence rates of best N -term Galerkin approximations for a class of elliptic SPDEs*, *Found. Comput. Math.*, 10 (2010), pp. 615–646.
- [8] N. COLLIER, A.-L. HAJI-ALI, F. NOBILE, E. VON SCHWERIN, AND R. TEMPONE, *A continuation multilevel Monte Carlo algorithm*, *BIT Numerical Mathematics*, 55 (2014), pp. 399–432.
- [9] P. R. CONRAD, Y. M. MARZOUK, N. S. PILLAI, AND A. SMITH, *Accelerating asymptotically exact MCMC for computationally intensive models via local approximations*, *J. Amer. Statist. Assoc.*, (2015).
- [10] G. M. CONSTANTINE AND T. H. SAVITS, *A multivariate Faà di Bruno formula with applications*, *Trans. Amer. Math. Soc.*, 248 (1996), pp. 503–520.
- [11] S. L. COTTER, G. O. ROBERTS, A. M. STUART, AND D. WHITE, *MCMC methods for functions: modifying old algorithms to make them faster*, *Stat. Sci.*, 28 (2013), pp. 424–446.
- [12] M. DASHTI AND A. M. STUART, *The Bayesian approach to inverse problems*, in *Handbook of Uncertainty Quantification*, R. Ghanem, D. Higdon, and H. Owhadi, eds., Springer, 2015.
- [13] W. DE AZEVEDO PRIBITKIN, *Simple upper bounds for partition functions*, *Ramanujan J.*, 18 (2009), pp. 113–119.
- [14] J. DICK, R. N. GANTNER, Q. T. LE GIA, AND CH. SCHWAB, *Higher order quasi-Monte Carlo integration for Bayesian estimation*, arXiv preprint arXiv:1602.07363, (2016).
- [15] J. DICK, R. N. GANTNER, Q. T. LE GIA, AND CH. SCHWAB, *Multilevel higher order quasi-Monte Carlo Bayesian estimation*, Tech. Report 2016-34, Seminar for Applied Mathematics, ETH Zürich, Switzerland, 2016.
- [16] J. DICK, F. Y. KUO, AND I. H. SLOAN, *High-dimensional integration: The quasi-Monte Carlo way*, vol. 22 of *Acta Num.*, Cambridge University Press, 2013, pp. 133–288.
- [17] C. R. DIETRICH AND G. N. NEWSAM, *Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix*, *SIAM J. Sci. Comput.*, 18 (1997), pp. 1088–1107.
- [18] T. J. DODWELL, C. KETELSEN, R. SCHEICHL, AND A. L. TECKENTRUP, *A hierarchical multilevel Markov chain Monte Carlo algorithm with applications to uncertainty quantification in subsurface flow*, *SIAM/ASA J. Uncertainty Quantification*, 3 (2015), pp. 1075–1108.
- [19] P. DOUKHAN AND G. LANG, *Evaluation for moments of a ratio with application to regression estimation*, *Bernoulli*, 15 (2009), pp. 1259–1286.
- [20] D. ELFVerson, F. HELLMAN, AND A. MÅLQVIST, *A multilevel Monte Carlo method for computing failure probabilities*, *SIAM/ASA Journal on Uncertainty Quantification*, 4 (2016), pp. 312–330.
- [21] E. C. FIELLER, *The distribution of the index in a normal bivariate population*, *Biometrika*, (1932), pp. 428–440.
- [22] C. GALEONE AND A. POLLASTRI, *Confidence intervals for the ratio of two means using the distribution of the quotient of two normals*, *Statistics in Transition*, 13 (2012), pp. 451–472.
- [23] R. N. GANTNER AND M. D. PETERS, *Higher order quasi-Monte Carlo for Bayesian shape inversion*, Tech. Report 2016-42, Seminar for Applied Mathematics, ETH Zürich, Switzerland, 2016.
- [24] R. G. GHANEM AND P. D. SPANOS, *Stochastic finite elements: a spectral approach*, Springer, New York, 1991.
- [25] M. B. GILES, *Multilevel Monte Carlo path simulation*, *Oper. Res.*, 256 (2008), pp. 981–986.
- [26] I. G. GRAHAM, F. Y. KUO, J. A. NICHOLS, R. SCHEICHL, CH. SCHWAB, AND I. H. SLOAN, *Quasi-Monte Carlo finite element methods for elliptic PDEs with lognormal random coefficients*, *Numer. Math.*, 131 (2015), pp. 329–368.

- [27] I. G. GRAHAM, F. Y. KUO, D. NUYENS, R. SCHEICHL, AND I. H. SLOAN, *Quasi-Monte Carlo methods for elliptic PDEs with random coefficients and applications*, J. Comput. Phys., 230 (2011), pp. 3668–3694.
- [28] ———, *Circulant embedding with QMC – analysis for elliptic PDE with lognormal coefficients*. In preparation, 2016.
- [29] A.-L. HAJI-ALI, F. NOBILE, E. VON SCHWERIN, AND R. TEMPONE, *Optimization of mesh hierarchies in multilevel Monte Carlo samplers*, Stochastics and Partial Differential Equations Analysis and Computations, 4 (2016), pp. 76–112.
- [30] S. HEINRICH, *Multilevel Monte Carlo methods*, vol. 2179 of Lecture notes in Comput. Sci., Springer, 2001, pp. 3624–3651.
- [31] D.V. HINKLEY, *On the ratio of two correlated normal random variables*, Biometrika, 56 (1969), pp. 635–639.
- [32] V. H. HOANG, CH. SCHWAB, AND A. M. STUART, *Complexity analysis of accelerated MCMC methods for Bayesian inversion*, Inverse Prob., 29 (2013), p. 085010.
- [33] C.A. KENNEDY AND W.C. LENNOX, *Moment operations on random variables, with applications for probabilistic analysis*, Probabilist. Eng. Mech., 16 (2001), pp. 253–259.
- [34] F. Y. KUO, R. SCHEICHL, CH. SCHWAB, I. H. SLOAN, AND E. ULLMANN, *Multilevel quasi-Monte Carlo methods for lognormal diffusion problems*, Preprint, arXiv:1507.01090 (2015).
- [35] F. Y. KUO, CH. SCHWAB, AND I. H. SLOAN, *Quasi-Monte Carlo finite element methods for a class of elliptic partial differential equations with random coefficients*, SIAM J. Numer. Anal., 50 (2012), pp. 3351–3374.
- [36] ———, *Multi-level quasi-Monte Carlo finite element methods for a class of elliptic PDEs with random coefficients*, Found. Comput. Math., 15 (2015), pp. 411–449.
- [37] F. Y. KUO AND I. H. SLOAN, *Lifting the curse of dimensionality*, Not. Am. Math. Soc., 52 (2005), pp. 1320–1328.
- [38] H. NIEDERREITER, *Random Number Generation and quasi-Monte Carlo methods*, SIAM, 1994.
- [39] D. ÖKSOY AND L. AROIAN, *The quotient of two correlated normal variables with applications*, Commun. Stat. Simulat., 23 (1994), pp. 223–241.
- [40] C. ROBERT AND G. CASELLA, *Monte Carlo Statistical Methods*, Springer, 1999.
- [41] C. SCHILLINGS AND C. SCHWAB, *Sparse, adaptive Smolyak quadratures for Bayesian inverse problems*, Inverse Probl., 29 (2013), p. 065011.
- [42] ———, *Scaling limits in computational Bayesian inversion*, SAM Research Report, ETH Zürich, 2014-26 (2014).
- [43] ———, *Sparsity in Bayesian inversion of parametric operator equations*, Inverse Probl., 30 (2014), p. 065007.
- [44] CH. SCHWAB AND A.M. STUART, *Sparse deterministic approximation of Bayesian inverse problems*, Inverse Probl., 28 (2012), p. 045003.
- [45] A. M. STUART, *Inverse problems: a Bayesian perspective*, vol. 19 of Acta Num., Cambridge University Press, 2010, pp. 451–559.
- [46] A. L. TECKENTRUP, *Multilevel Monte Carlo methods and uncertainty quantification*, PhD thesis, University of Bath, 2013.

- [47] A. L. TECKENTRUP, R. SCHEICHL, M. B. GILES, AND E. ULLMANN, *Further analysis of multilevel Monte Carlo methods for elliptic PDEs with random coefficients*, Numer. Math., 125 (2013), pp. 569–600.
- [48] H. TRAN, C.G. WEBSTER, AND G. ZHANG, *Analysis of quasi-optimal polynomial approximations for parameterized PDEs with deterministic and stochastic coefficients*, arXiv preprint arXiv:1508.01821, (2015).