Certified Reduced Basis Model Characterization: a Frequentistic Uncertainty Framework[☆]

D.B.P. Huynh[a], D.J. Knezevic[a], A.T. Patera[a,*]

[a] *Department of Mechanical Engineering and Center for Computational Engineering Massachusetts Institute of Technology, Cambridge, MA, 02139 USA*

Abstract

We introduce a frequentistic validation framework for assessment — acceptance or rejection — of the consistency of a proposed parametrized partial differential equation model with respect to (noisy) experimental data from a physical system. Our method builds upon the Hotelling T^2 statistical hypothesis test for bias first introduced by Balci & Sargent in 1984 and subsequently extended by McFarland & Mahadevan 2008. Our approach introduces two new elements: a spectral representation of the misfit which reduces the dimensionality and variance of the underlying multivariate Gaussian but without introduction of the usual regression assumptions; a certified (verified) reduced basis approximation — reduced order model — which greatly accelerates computational performance but without any loss of rigor relative to the full (finite element) discretization. We illustrate our approach with examples from heat transfer and acoustics, both based on synthetic data. We demonstrate that we can efficiently identify possibility regions that characterize parameter uncertainty; furthermore, in the case that the possibility region is empty, we can deduce the presence of "unmodeled physics" such as cracks or heterogeneities.

Keywords: Validation; uncertainty quantification; hypothesis testing; partial differential equation; reduced basis method; *a posteriori error* bound

1. Introduction

Parameterized models play a central role in engineering. Our focus here is parametrized partial differential equations (PDEs) that arise for example in heat transfer, solid mechanics, acoustics, fluid flow, and electromagnetics. The parameters appear as coefficients in the partial differential equation and may represent physical properties, interactions with the environment, or geometry. Often one set of parameters will serve

[☆]Submitted to CMAME.

*Corresponding author.

[1]Email addresses: (D.B.P. Huynh) huynh@mit.edu, (D.J. Knezevic) dknez@mit.edu, (A.T. Patera) patera@mit.edu.

to relate the PDE model to the actual Physical System (PS) of interest, while another set of parameters will then serve to subsequently design or optimize or control the Physical System. Our interest in this paper is the former.

There are many frameworks by which to relate a proposed PDE model to the Physical System of interest. *Parameter estimation* [1, 2] and *calibration* (or "update") [3, 4, 5] search for a best choice of parameters relative to (typically noisy) experimental data. *Validation* [3] assesses the consistency of any candidate parameter value, and hence associated proposed PDE model, with (typically noisy) experimental data. Both activities — calibration and in particular validation — play a central role in Uncertainty Quantification ("UQ"). Our focus in this paper is on validation.

Validation procedures fall into two broad categories: Bayesian and frequentistic [6]. Bayesian validation builds directly upon Bayesian parameter estimation [2, 4] — from experimental data to likelihood through prior to posterior and finally credible regions [6]. The Bayesian approach is quite general and can yield actionable information on many quantities of interest; however, the prior can be difficult to specify or interpret [3, 7], and hence absolute conclusions can remain elusive. Frequentistic validation procedures [8, 7] are typically constructed as statistical tests of appropriate hypotheses on bias which then permit acceptance or rejection of a particular choice of candidate parameters. The frequentistic approach avoids the Bayesian prior and can thus provide more rigorous or at least less subjective conclusions; however, Type II errors [9] can typically not be controlled, and hence the implications are perforce restricted.

In this paper we consider frequentistic validation. The seminal work in frequentistic validation, due to Balci and Sargent [8], is based on the multivariate Gaussian Hotelling T^2 test [10]. More recently, McFarland and Mahadevan [7] extend and apply the Balci & Sargent approach to the thermal validation challenge problem proposed by Sandia National Laboratory [11]. We shall consider in this paper two further innovations to the methods proposed in [8, 7]: the first, (arguably) an improvement, is introduced to reduce the dimensionality of multivariate Gaussian; the second, an extension, is introduced to permit more efficient computation. We discuss each in turn.

As regards the first innovation, we develop the hypothesis test based not on pointwise data but rather on "spectral" Legendre coefficients [12]. This approach has two advantages: it leads to a significant reduction in the dimensionality of the multivariate Gaussian and hence permits more rigorous inference in particular of the variance (e.g., without further hypotheses); it permits inferences from fewer experimental realizations (or replications) since the spectral coefficients implicitly average the data and hence reduce the variance, at least for correlation scales small compared to the full observation window.

We emphasize that our approach differs from standard regression [13] in that we do *not* assume that our truncated spectral expansion is unbiased: our inferences derive from a Linear Optimization problem constrained by the Legendre coefficient joint confidence intervals. Note also that we do not assume uncorrelated noise and in fact we permit any Gaussian process [14]: of course, unlike the standard regression framework, we can not simultaneously deduce bias and variance from a single realization; nevertheless, as already indicated, relatively few replications suffice.

Validation perforce requires the consideration of many candidate parameter values. In the PDE context, the evaluation of the model by classical (in our case) finite element [15] methods can be very expensive. Our second innovation is thus to replace the "truth" finite element discretization with a less costly surrogate: a certified reduced basis approximation [16, 17, 18, 19, 20, 21] which provides not just an inexpensive output prediction but also — in the sense of "verification" — an inexpensive but rigorous output error bound. The reduced basis (RB) approach is efficient in the many–query and real–time contexts, both of which are relevant in (calibration and) validation procedures. In order to incorporate the certified reduced basis method into the frequentistic validation framework we consider a composite hypothesis [9, 8]. The RB error bound is crucial in ensuring proper distinction between experimental noise, bias in the model prediction, and numerical (reduced basis) error: without the RB error bound, we might conflate numerical error with bias and hence reject parameter values which in fact are consistent with experimental data. (We note that metamodels and reduced models are gainfully invoked [22, 7, 5, 23] in several calibration and validation approaches, but only in a few studies [24] are rigorous error bounds considered.)

The rest of this paper is arranged as follows. In Section 2 we introduce the statistical framework — a hypothesis test for bias — for our approach independent of any specific mathematical model or physical system. In Section 3 we specialize the framework to the case in which our model is given by a proposed parametrized PDE; in particular, we incorporate the reduced basis output *and* reduced basis output error bound into our hypothesis test. In Section 4, we present numerical results for a transient thermal conduction problem and a Helmholtz acoustics problem. In both cases we demonstrate the ability of our approach to identify "possibility regions" based on many-query parameter sweeps over the parameter domain, and also to deduce the presence of "unmodeled physics" (defects) in a Physical System.

2. A Statistical Result

We introduce a compact domain $\mathcal{D}_\nu \in \mathbb{R}^{P_\nu}$; in this paper we shall exclusively consider $P_\nu = 1$, and hence we may write $\mathcal{D}_\nu \equiv [\nu_{\min}, \nu_{\max}]$. We then define a (centered) Gaussian random process [14] $G(\nu, \omega)$ such

that the mean is zero,

$$\int d\mathbb{P}(\omega) G(\nu, \omega) = 0, \forall \nu \in \mathcal{D}_\nu, \tag{1}$$

and the average variance is unity,

$$\frac{1}{\nu_{\max} - \nu_{\min}} \int d\mathbb{P}(\omega) \int_{\mathcal{D}_\nu} G^2(\nu, \omega) d\nu = 1. \tag{2}$$

Finally, for given positive definite covariance kernel C, G satisfies

$$\int d\mathbb{P}(\omega) G(\nu, \omega) G(\nu', \omega) = C(\nu, \nu'), \quad \forall \nu, \nu' \in \mathcal{D}_\nu. \tag{3}$$

Here $d\mathbb{P}(\omega)$ indicates integration with respect to the probability measure. Note that the variance condition on G, (2), represents a scaling constraint on the covariance C.

In actual practice such a Gaussian process can be represented by a Karhunen-Loève expansion [25]. We first find $\chi : \mathcal{D}_\nu \to \mathbb{R}$ and $\lambda \in \mathbb{R}$ solutions of the eigenproblem

$$\int_{\mathcal{D}_\nu} C(\nu, \nu') \chi(\nu') d\nu' = \lambda \chi(\nu) \tag{4}$$

with normalization $\int_{\mathcal{D}_\nu} \chi^2(\nu) d\nu / (\nu_{\max} - \nu_{\min}) = 1$. We may then express our random process as

$$G(\nu, \omega) = \sum_{l=1}^{\infty} \sqrt{\lambda_l} W_l(\omega) \chi_l(\nu), \tag{5}$$

where the W_l are independent normally distributed random variables with zero mean and unity variance. Note we order the (positive) eigenvalues in decreasing magnitude ($\lambda_1 \geq \lambda_2 \cdots$). The sum of the eigenvalues is the integrated variance, and thus from (2) we conclude that

$$\frac{1}{\nu_{\max} - \nu_{\min}} \sum_{l=1}^{\infty} \lambda_l = 1.$$

In practice we truncate the expansion (5) at L terms such that sum of the remaining eigenvalues $\lambda_{L+1}, \lambda_{L+2}, \ldots$ is small compared to unity.

In our procedure we shall assume that the underlying noise or stochastic element is a Gaussian process. However — and unlike kriging [26] — we will *not* need to know the details of the process in order to construct our approximation or provide our statistical inferences. Nevertheless, for purposes of generating

representative synthetic data, we will consider two convenient correlation functions:

(i) White noise:

$$C(\nu,\nu') = \begin{cases} \text{Const}, & \nu = \nu' \\ 0, & \nu \neq \nu' \end{cases} \tag{6}$$

(ii) Correlated noise:

$$C(\nu,\nu') = \text{Const}\exp(-|\nu-\nu'|/\Lambda_\nu), \tag{7}$$

where Λ_ν is a specified correlation scale in ν.

In each case Const is chosen to ensure unity average variance over $\mathcal{D}_\nu$, (2). With increasing correlation scale the spectrum of (4) decreases more rapidly and we may choose L correspondingly smaller — we may include fewer and fewer terms in the truncated Karhunen–Loève expansion.

We now introduce a function $\delta : \mathcal{D}_\nu \to \mathbb{R}$. We then define, for "given" integrated variance ϵ, the random process

$$Y(\nu,\omega) \equiv \delta(\nu) + \epsilon G(\nu,\omega). \tag{8}$$

Note that the expectation of $Y(\nu,\cdot)$ is $\delta(\nu)$ from our assumption (1) on G. We emphasize that in our inference procedure we will *not* need to know the magnitude of the average (over $\mathcal{D}_\nu$) variance, ϵ.

We next introduce a set of "quadrature points" (ultimately measurement points) $\nu_m \in \mathcal{D}_\nu, 1 \leq m \leq M$, and associated positive quadrature weights $\rho_m, 1 \leq m \leq M$. We then provide a "projection matrix"

$$w_m^{[i]} = \frac{\rho_m}{\nu_{\max}-\nu_{\min}} L_i\left(-1+2\left(\frac{\nu_m-\nu_{\min}}{\nu_{\max}-\nu_{\min}}\right)\right), \quad 1 \leq m \leq M,\ 0 \leq i \leq I, \tag{9}$$

where $L_i : [-1,1] \to \mathbb{R}$ is the Legendre polynomial of order i [27]. The "Legendre coefficients" of δ are then given by

$$\delta^{[i]} = \sum_{m=1}^{M} w_m^{[i]} \delta(\nu_m), \quad 0 \leq i \leq I. \tag{10}$$

We note that if δ is a smooth function then we expect rapid decrease in the Legendre coefficients.

There is much freedom in the choice of quadrature (points and hence weights). In this paper we employ the trapezoidal rule:

$$\nu_m = \nu_{\min} + (m-1)\Delta\nu, \quad \rho_m = c_m \Delta\nu, \quad 1 \leq m \leq M, \tag{11}$$

where $\Delta\nu = (\nu_{\max}-\nu_{\min})/(M-1)$, and $c_1 = c_M = 1/2$, $c_m = 1$ for $2 \leq m \leq M-1$. Clearly there may be some advantage to a Gauss–Legendre (or Gauss–Lobatto) quadrature [27], however the associated clustered

point distribution could be unnatural in an experimental implementation and we thus restrict attention to low–order schemes.

Now let

$$Y_m^{\rm exp}(\omega) \equiv Y(\nu_m, \omega), \quad 1 \le m \le M, \tag{12}$$

be a realization ω of our random process (8): a particular set of experimental measurements of $Y(\nu, \omega)$ at the points $\nu_m, 1 \le m \le M$. We may then introduce

$$Z^{[i]}(\omega) = \sum_{m=1}^{M} w_m^{[i]} Y_m^{\rm exp}(\omega), \quad 0 \le i \le I; \tag{13}$$

we may view $Z^{[i]}(\omega)$ as a (noisy) approximation to the Legendre coefficient $\delta^{[i]}$ of (10). It is instructive to rewrite (13) in terms of our Karhunen–Loève expansion (now truncated at L terms): from (5), (8), and (13) we obtain

$$Z^{[i]}(\omega) = \delta^{[i]} + \epsilon \sum_{l=1}^{L} W_l(\omega) [\sqrt{\lambda_l} \sum_{m=1}^{M} w_m^{[i]} \chi_l(\nu_m)], \quad 0 \le i \le I. \tag{14}$$

It then follows that the Legendre coefficient random vector, $(Z^{[0]}, Z^{[1]}, \ldots, Z^{[I]})(\omega)$, is $(I+1)$–variate normally distributed with mean $(\delta^{[0]}, \delta^{[1]}, \ldots, \delta^{[I]})$ [9].

We next introduce an ensemble of K realizations (presumed of course independent) $\omega_1, \ldots, \omega_K$: in practice this ensemble will take the form of K sets of M measurements each. We may then introduce the sample mean of the Legendre coeffients over this ensemble as

$$\overline{Z}_K^{[i]} = \frac{1}{K} \sum_{k=1}^{K} Z^{[i]}(\omega_k), \quad 0 \le i \le I, \tag{15}$$

for $Z^{[i]}(\omega)$ given by (13). Similarly, the sample variance over the ensemble is given by

$$\mathcal{S}_K^{[i]} = \frac{1}{K-1} \sum_{k=1}^{K} (Z^{[i]}(\omega_k) - \overline{Z}_K^{[i]})^2, \quad 0 \le i \le I. \tag{16}$$

(We shall not need the full sample covariance for reasons described below.)

It then follows that the intervals

$$\mathcal{I}_K^{[i]} \equiv \left[\overline{Z}_K^{[i]} - T_\gamma \sqrt{\frac{\mathcal{S}_K^{[i]}}{K}}, \overline{Z}_K^{[i]} + T_\gamma \sqrt{\frac{\mathcal{S}_K^{[i]}}{K}} \right], \quad 0 \le i \le I, \tag{17}$$

represent *joint* $100\gamma\%$ confidence intervals for the $\delta^{[i]}, 0 \leq i \leq I$:

$$\Pr\{\delta^{[i]} \in \mathcal{I}_K^{[i]}, 0 \leq i \leq I\} \geq \gamma. \tag{18}$$

Here

$$T_\gamma \equiv \sqrt{\frac{(K-1)(I+1)}{K-I-1} F(\gamma, I+1, K-I-1)}, \tag{19}$$

where $F(\gamma, \cdot, \cdot)$ denotes the γ quantile of the F distribution.

We make four points concerning this inference. First, we note that

$$\mathcal{I}_K^{[0]} \times \mathcal{I}_K^{[1]} \times \cdots \mathcal{I}_K^{[I]}$$

represents a rectangular confidence region which encloses the strict elliptical confidence region [10] — the i^{th} dimension of this rectangular region is derived by maximizing the i^{th} component of $y \in \mathbb{R}^{I+1}$, for $i = 1, \ldots, I+1$, with y constrained to reside on the elliptical confidence region from [10]. We prefer the rectangular region in anticipation of the Linear Program introduced below; ultimately, however, the elliptical region (or an oriented rectangular region) could be incorporated and would of course sharpen our results. Second, we emphasize that our inference related to the $I+1$ Legendre coefficients does not assume — in contrast to a corresponding $(I+1)$–term regression model [13] — that the bias is zero: our inference (18) does *not* suppose that $\delta(\nu)$ is exactly represented by the first $(I+1)$ terms of a Legendre expansion. Third, we observe that typically the number of measurements M might be quite large compared to the requisite number of Legendre coefficients $I+1$; this will, in view of the denominator in (19), permit inference with fewer replications than in pointwise approaches (e.g., McFarland & Mahadevan [7]). Finally, we implicitly obtain variance reduction since the $Z^{[i]}$ are in some sense "moments" constructed as averages over the M measurements; this further reduces the number of replications required.

We elaborate further on the variance reduction. We recall that the procedure defined above requires an ensemble of K realizations (which we may also view as repetitions or replications) each of which comprises M measurements. In general, given that we allow correlated errors — often the case, for example if ν represents time or space and disturbances and measurement errors are spectrally broad — a single realization can not suffice. As a simple counterexample we might consider the case in which G has infinite correlation scale: a single realization can not possibly distinguish signal from noise. However, our procedure does nevertheless benefit from the M data points associated with each realization: if G has relatively short correlation scale

Λ_ν — in the limit $\Lambda_\nu \to 0$ we recover the standard regression assumption [13] of white Gaussian noise — then the variance of the pseudo sample means $Z^{[i]}(\omega), 0 \le i \le I$, will be small and only very few realizations will be required in order to obtain an acceptably tight confidence interval.

We now define

$$\delta_{\max} \equiv \max_{m \in \{1,\ldots,M\}} |\delta(\nu_m)|. \tag{20}$$

We shall then wish to test the hypothesis

$$\mathcal{H} \equiv \{\ \delta_{\max} \le B\ \} \tag{21}$$

for some positive constant B. We may view (21) as a hypothesis on bias; we introduce in the next section the particular interpretation of bias relevant in our context.

We now describe the rejection criterion. We shall reject hypothesis $\mathcal{H}$ if and only if

$$\hat{\delta}_{\max} > B, \tag{22}$$

where

$$\hat{\delta}_{\max} = \min_{v \in \mathcal{R}} \max_{m \in \{1,\ldots,M\}} |v_m|, \tag{23}$$

and the constraint set $\mathcal{R}$ is given by

$$\mathcal{R} = \{v \in \mathbb{R}^M \mid \sum_{m=1}^{M} w_m^{[i]} v_m \in \mathcal{I}_K^{[i]}, 0 \le i \le I\}. \tag{24}$$

We shall denote by $1-\gamma$ (where γ here determines T_γ in (19)) the size of the test.

In particular, we now derive a bound for the Type I [9] error in

Proposition 2.1. *The probability of rejection of hypothesis $\mathcal{H}$ when hypothesis $\mathcal{H}$ is in fact true is less than or equal to $1-\gamma$.*

Proof. We consider, as dictated by our hypothesis, the case in which the hypothesis $\mathcal{H}$ is rejected: $\hat{\delta}_{\max} > B$. We first suppose that the condition $\mathcal{E} \equiv \{\delta^{[i]} \in \mathcal{I}_K^{[i]}, 0 \le i \le I\}$ is satisfied, which imples that $\delta(\nu_m), 1 \le m \le M$, is in the admissible set $\mathcal{R}$. Since $\hat{\delta}_{\max}$ is defined in (23), (24), as the *minimum* over $v \in \mathcal{R}$ of $\max_{m \in \{1,\ldots,M\}} |v_m|$, it follows that $\delta_{\max} \ge \hat{\delta}_{\max}$.

Then, under the above conditions, we may conclude that $\delta_{\max} > B$, which implies that the hypothesis

$\mathcal{H}$ is indeed false. Hence $\mathcal{E}$ true implies correct rejection of $\mathcal{H}$, and thus a necessary condition for incorrect rejection of hypothesis $\mathcal{H}$ (Type I error) is $\mathcal{E}$ false. We conclude from (18) that the probability of a Type I error is less than or equal to $1 - \gamma$. □

Remark 2.1. *We can not control the Type II error [9] — in which we accept hypothesis $\mathcal{H}$ when hypothesis $\mathcal{H}$ is in fact false — though we do anticipate that for larger K (tighter confidence intervals) and larger I (better estimates for $\delta_{\max}$) the probability of Type II errors shall decrease. Similarly, we are not able to quantify the power of our test.*

Finally, we note that our test (22)–(24) can be efficiently computed. It is a standard procedure to reformulate (23), (24), as a simple Linear Program (LP) in $M+1$ variables and $2(M+I+1)$ constraints as follows: $\hat{\delta}_{\max} \equiv \min c$ over $(c, v_1, \ldots, v_M) \in \mathbb{R}^{M+1}$, where $(c, v_1, \ldots, v_M) \in \mathbb{R}^{M+1}$ satisfies

$$\begin{aligned} -c \le v_m \le c, \qquad & m = 1, \ldots, M, \\ \mathcal{I}^{[i]}_{K,\min} \le \sum_{m=1}^{M} w^{[i]}_m v_m \le \mathcal{I}^{[i]}_{K,\max}, \quad & i = 0, \ldots, I. \end{aligned}$$

Note each two-sided inequality represents two constraints.

3. PDE-Based Statistical Inference

We now proceed to the parametrized partial differential equation context. In addition to $\nu \in \mathcal{D}_\nu$ which is a control parameter (or set of parameters) we also introduce a parameter $\kappa \in \mathcal{D}_\kappa \subset \mathbb{R}^{P_\kappa}$ (ultimately, we shall wish to determine κ from data). We are given a *proposed* parametrized partial differential equation model $\mathcal{M}^{\nu,\kappa}$; this partial differential equation yields an exact (model–prediction) output $s\colon \mathcal{D}_\nu \times \mathcal{D}_\kappa \to \mathbb{R}$. In actual practice s may not be obtained analytically, and instead we resort to a high-fidelity finite element [15] "truth" approximation $\mathcal{M}^{\nu,\kappa}_h$ (here h denotes the mesh size); this approximation yields a truth output $s_h : \mathcal{D}_\nu \times \mathcal{D}_\kappa \to \mathbb{R}$. We presume that the truth approximation is sufficiently rich — sufficiently many and judiciously chosen degrees of freedom (h sufficiently small) — such that the difference between s and s_h is negligible. Unfortunately, this truth approximation will also often be prohibitively costly: each evaluation $(\nu, \kappa) \to s_h(\nu, \kappa)$ will be very expensive.

We thus introduce a certified reduced basis approximation [20, 28, 29, 30] which takes advantage of the parametric definition and parametric structure of the problem to greatly reduce the computational cost in the many–query and real–time contexts. We shall denote the N degree–of–freedom reduced basis approximation

by $\mathcal{M}_N^{\nu,\kappa}$; this reduced basis approximation yields a reduced basis output $s_N : \mathcal{D}_\nu \times \mathcal{D}_\kappa \to \mathbb{R}$ *and* a reduced basis output error bound $\Delta_N : \mathcal{D}_\nu \times \mathcal{D}_\kappa \to \mathbb{R}$ such that $|s_h(\nu,\kappa) - s_N(\nu,\kappa)| \leq \Delta_N(\nu,\kappa), \forall (\nu,\kappa) \in \mathcal{D}_\nu \times \mathcal{D}_\kappa$. The reduced basis method accommodates an Offline–Online computational strategy which accepts "many–truth" cost in the Offline stage in exchange for very rapid evaluation $(\nu,\kappa) \to s_N(\nu,\kappa), \Delta_N(\nu,\kappa)$ in the Online (many–query, real–time) stage. Note that in actual practice we consider a hierarchy of reduced basis approximations $\mathcal{M}_N^{\nu,\kappa}$, $1 \leq N \leq N_{\max}$.

We assume that our experimental output is given by the random Gaussian process

$$V(\nu,\omega) \equiv s^{\mathrm{PS}}(\nu) + \epsilon G(\nu,\omega) \tag{25}$$

where $s^{\mathrm{PS}}(\nu)$ represents the response of our physical system. Note that ϵ and G are *unknown to us*: we presume only that the random process G is Gaussian.

We now wish to compare the output of our proposed parametrized PDE model $\mathcal{M}_h^{\nu,\kappa}$ to the measured data $V(\nu,\omega)$. We thus form the new random variable $D(\nu,\kappa,\omega) \equiv V(\nu,\omega) - s_N(\nu,\kappa)$ or

$$D(\nu,\kappa,\omega) \equiv s^{\mathrm{PS}}(\nu) - s_N(\nu,\kappa) + \epsilon G(\nu,\omega). \tag{26}$$

We then presume that we are given some candidate value $\kappa = \hat{\kappa}$ in $\mathcal{D}_\kappa$ and define $Y^{\hat{\kappa}}(\nu,\omega) \equiv D(\nu,\hat{\kappa})$ or

$$Y^{\hat{\kappa}}(\nu,\omega) \equiv [s^{\mathrm{PS}}(\nu) - s_N(\nu,\hat{\kappa})] + \epsilon G(\nu,\omega). \tag{27}$$

We note that (27) has the same form as (8) but now $\delta(\nu) \equiv \delta^{\hat{\kappa}}(\nu) \equiv s^{\mathrm{PS}}(\nu) - s_N(\nu,\hat{\kappa})$ has the particular interpretation of the deviation or misfit between experiment and (reduced basis) prediction.

We now proceed to choose the number of measurements per realization, M, the associated measurement and quadrature points $\nu_m, \rho_m, 1 \leq m \leq M$, the order of the Legendre approximation I, and finally the sample size K. We then perform the necessary experiments in order to evaluate the $Z^{[i]}(\omega_k), 1 \leq k \leq K, 0 \leq i \leq I$, of (13), and subsequently the $\overline{Z}_K^{[i]}, 0 \leq i \leq I$, of (15). Note that in this context (12) should be interpreted (for given $\hat{\kappa}$) as $Y_m^{\exp}(\omega) \equiv V^{\exp}(\nu_m,\omega) - s_N(\nu_m,\hat{\kappa})$: for $1 \leq m \leq M$, we measure $V^{\exp}(\nu_m,\omega) = s^{\mathrm{PS}}(\nu_m) + \epsilon G(\nu_m,\omega)$ and then subtract the reduced basis prediction $s_N(\nu_m,\hat{\kappa})$.

Finally, we now set

$$B^{\hat{\kappa}} = \max_{m \in \{1,\ldots,M\}} \Delta_N(\nu_m,\hat{\kappa}) \tag{28}$$

and we consider the hypothesis (21) with $B \equiv B^{\hat{\kappa}}$: the bound in our test is now related to the maximum

reduced basis error bound over the measurements points $\nu_m, m = 1, \ldots, M$. Note in this context our hypothesis reads

$$\mathcal{H}^{\hat{\kappa}} \equiv \{|\delta^{\hat{\kappa}}(\nu_m)| \leq B^{\hat{\kappa}}, 1 \leq m \leq M\}. \tag{29}$$

We then perform the statistical test of Section 2 for the selected confidence level γ.

If *we accept the hypothesis* we know that we can explain the difference between the experimental measurements and the reduced basis prediction in terms of the reduced basis approximation error: the proposal $\hat{\kappa}$ should thus be included as a viable (or consistent) candidate for κ^*. We recall that we can not conclude that the truth model is indeed valid simply because we do not reject a value $\hat{\kappa}$: we do not control the Type II error. On the other hand, if *we reject the hypothesis* then we know with confidence greater than γ that the proposed *truth* model $\mathcal{M}_h^{\nu,\hat{\kappa}}$ is not consistent with the experimental data. In particular, the discrepancy between experiment and model is larger than can be attributed to the error incurred in the reduced basis approximation of the truth:

$$\max_{\nu_m} |s^{\mathrm{PS}}(\nu_m) - s_h(\nu_m, \hat{\kappa})| + \max_{\nu_m} |s_h(\nu_m, \hat{\kappa}) - s_N(\nu_m, \hat{\kappa})| \geq \max_{\nu_m} |s^{\mathrm{PS}}(\nu_m) - s_N(\nu_m, \hat{\kappa})| > \max_{\nu_m} \Delta_N(\nu_m, \hat{\kappa}),$$

and hence since $\max_{\nu_m} |s_h(\nu_m, \hat{\kappa}) - s_N(\nu_m, \hat{\kappa})| \leq \max_{\nu_m} \Delta_N(\nu_m, \hat{\kappa})$ it is not possible that $s^{\mathrm{PS}}(\nu) = s_h(\nu, \hat{\kappa})$. The cause for the discrepancy must thus be *either* a proposed value $\hat{\kappa}$ which differs from the actual physical value κ^* *or* some physical effects that are simply not included in our proposed PDE model. Note that our hypothesis can not distinguish between these two possibilities.

We may now consider many values of $\hat{\kappa}$. (In actual practice, we would presumably only search in some neighborhood of the (or a) nonlinear least–squares "misfit" minimizer provided by (in our implementations) the Levenberg–Marquardt algorithm [31].) The set

$$\mathcal{P}_\kappa \equiv \{\hat{\kappa} \in \mathcal{D}_\kappa \,|\, \mathcal{H}^{\hat{\kappa}} \text{ accepted } \}$$

constitutes a parameter uncertainty region, or more optimistically a parameter "possibility" region: $\mathcal{P}_\kappa$ includes all values of $\hat{\kappa}$ for which the discrepancy between data and model prediction is consistent with the reduced basis accuracy. Note that $\mathcal{P}_\kappa$ will depend on the accuracy of the reduced basis approximation, the values of γ, K, M and I, and of course also the sensitivity of the output to the parameters. The test of hypothesis $\mathcal{H}^{\hat{\kappa}}$ for each $\hat{\kappa}$ requires evaluation of $s_N(\nu_m, \hat{\kappa}), 1 \leq m \leq M$: these calculations can be readily effected thanks to the extremely low cost of the Online reduced basis output evaluation. In essence,

the reduced basis *output approximation* permits us to address the parameter uncertainty issue — direct calculation of the truth output $s_h(\nu_m, \hat{\kappa}), 1 \leq m \leq M$, for many $\hat{\kappa}$ in $\mathcal{D}_\kappa$ would be prohibitive (certainly in real–time); the reduced basis *error bound*, in conjunction with our hypothesis test, permits us to nevertheless provide rigorous conclusions with respect to the *truth* approximation.

Of notable interest is the case in which the parameter possibility region is empty: for all $\hat{\kappa} \in \mathcal{D}_\kappa$ we reject our hypothesis. (Note we can no longer attribute a confidence level to this "joint" outcome since our Type I error analysis considers only a single value of $\hat{\kappa}$.) In this case we would conclude that the misfit can not be attributed to reduced basis error and hence the physical system of interest can *not* be adequately represented by the proposed truth PDE model — we suppose here that $\mathcal{D}_\kappa$ is sufficiently large (or includes all plausible parameter values) such that we can safely attribute the rejections to unmodeled physics. This situation is most instructive if we can interpret the unmodeled physics as *changes* to a system such as the quasi–static appearance of defects or geometric variations relative to an initial baseline, or perhaps manufacturing departures from nominal specifications; if our hypothesis is rejected for all $\hat{\kappa}$ then we can conclude that the system is indeed appreciably different from our reference or desired configuation. This "detection" capability can be very effective as it permits us to exploit a relatively low–dimensional reduced basis parametrization — $\mathcal{D}_\nu \times \mathcal{D}_\kappa$ — to deduce potentially very high–dimensional (and important) deviations — for example, the appearance of defects (e.g. voids or cracks or tumors) in a physical system.

Finally, we could also include a perturbative term on the right–hand side of our proposed partial differential equation to account for unmodeled physics. We may not know the specific details of the perturbation (for example, spatial fluctuations in the thermal conductivity on a fine scale), but based on plausible assumptions and our reduced basis stability factors we may be able to bound the effect of the perturbation on the output. We can then include this additional "allowable deviation" in our choice of B in (21): now B will be the sum of the reduced basis error and unmodeled physics contributions. We do not consider this extension in the numerical results section here, but it is a possible direction for more robust treatment of (inevitable) unmodeled physics in future work.

4. Numerical Results

We illustrate the frequentistic uncertainty framework developed above for two example problems. The first example is a transient thermal conduction problem with uncertain thermal conductivity parameters; the second example is a Helmholtz acoustics problem. A `libMesh`-based implementation of the certified Reduced Basis method [32, 33] is employed to generate our finite element and RB numerical results.

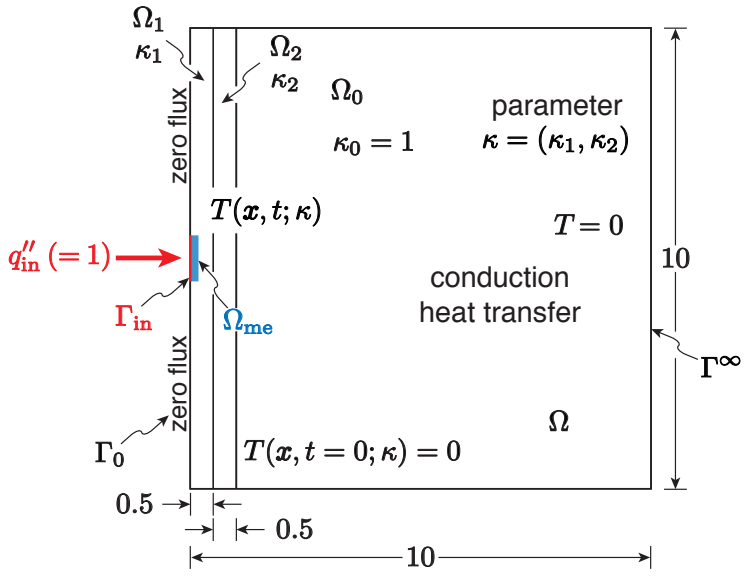

Figure 1: The transient thermal conduction problem. The origin of the $\boldsymbol{x} = (x_1, x_2)$ coordinate system is the lower left corner.

4.1. Transient Thermal Conduction

Proposed PDE Model

We consider transient thermal conduction in the domain depicted in Figure 1. We consider two layers of material Ω_1 and Ω_2 with respective parametrized conductivities κ_1 and κ_2; these two layers are affixed to a substrate Ω_0 with specified (normalized) conductivity $\kappa_0 \equiv 1$. We impose a heat flux on the surface and we take as our output the average temperature in an adjacent "measurement region" as a function of time. We now present the the mathematical statement of the proposed parametrized PDE model. Note that the proposed parametrized PDE model shall include both the equation for the field as well as the prescription for the output.

We consider the computational domain $\Omega \equiv \Omega_0 \cup \Omega_1 \cup \Omega_2$, where $\Omega_0 \equiv (2, 10) \times (0, 10)$, $\Omega_1 \equiv (0, 1) \times (0, 10)$, and $\Omega_2 \equiv (1, 2) \times (0, 10)$. We define $\Gamma_{\text{in}} \equiv \{x \in \partial\Omega : x_1 = 0 \text{ and } x_2 \in (4.5, 5.5)\}$, $\Gamma^{\infty} \equiv \{x \in \partial\Omega : x_2 = 0 \text{ or } x_2 = 10 \text{ or } x_1 = 10\}$, and the Sobolev space $X \equiv \{v \in H^1(\Omega) : v = 0 \text{ on } \Gamma^{\infty}\}$. We consider the time interval $t \in [0, t_f]$, where $t_f = 5$; we divide the time interval into $J = 100$ subintervals of equal length $\Delta t = 0.05$ and let $t^j \equiv j\Delta t$, $0 \le j \le J$. We consider the two-dimensional parameter domain $\kappa \equiv (\kappa_1, \kappa_2) \in \mathcal{D}_\kappa \equiv [0.2, 5]^2$. Note for this problem ν from Section 2 corresponds to time (hence $\nu_{\min} = 0$ and $\nu_{\max} = t_f$) and κ from Section 3 corresponds, in fact, to κ.

We first consider $\mathcal{M}^{\nu,\kappa}$, the (semidiscrete backward Euler) weak form of the governing PDE: For $\kappa \equiv$

$(\kappa_1, \kappa_2) \in \mathcal{D}_\kappa$ and $j = 1, \ldots, J$, find $u(t^j, \kappa) \in X$ that satisfies

$$\int_\Omega \frac{u(t^j,\kappa) - u(t^{j-1},\kappa)}{\Delta t} v + \int_{\Omega_0} \nabla u(t^j,\kappa) \cdot \nabla v + \sum_{i=1}^{2} \int_{\Omega_i} \kappa_i \nabla u(t^j,\kappa) \cdot \nabla v = \int_{\Gamma_{\text{in}}} v, \quad \forall v \in X. \tag{30}$$

We also introduce the output functional

$$\ell(v) = \frac{1}{|\Omega_{\text{me}}|} \int_{\Omega_{\text{me}}} v, \quad \forall v \in X, \tag{31}$$

where $\Omega_{\text{me}} \equiv (0, 0.2) \times (4.5, 5.5)$; our output is then given by $s(t^j, \kappa) = \ell(u(t^j, \kappa))$ for $0 \le j \le J$.

To obtain $\mathcal{M}_h^{\nu,\kappa}$ we now replace X in (30) with a truth finite element space $X_h \subset X$ with $\dim(X_h) = 3009$ degrees of freedom: For $\kappa \equiv (\kappa_1, \kappa_2) \in \mathcal{D}_\kappa$ and $j = 1, \ldots, J$, $u_h(t^j, \kappa) \in X_h$ satisfies

$$\int_\Omega \frac{u_h(t^j,\kappa) - u_h(t^{j-1},\kappa)}{\Delta t} v + \int_{\Omega_0} \nabla u_h(t^j,\kappa) \cdot \nabla v + \sum_{i=1}^{2} \int_{\Omega_i} \kappa_i \nabla u_h(t^j,\kappa) \cdot \nabla v = \int_{\Gamma_{\text{in}}} v, \quad \forall v \in X_h. \tag{32}$$

The truth output is then given by $s_h(t^j, \kappa) = \ell(u_h(t^j, \kappa))$ for $0 \le j \le J$.

The RB approximation $u_N(t^j, \kappa)$ also inherits the Galerkin formulation from (30), though now the approximation space is the RB space X_N: For $\kappa \equiv (\kappa_1, \kappa_2) \in \mathcal{D}_\kappa$ and $j = 1, \ldots, J$, $u_N(t^j, \kappa) \in X_N$ satisfies

$$\int_\Omega \frac{u_N(t^j,\kappa) - u_N(t^{j-1},\kappa)}{\Delta t} v + \int_{\Omega_0} \nabla u_N(t^j,\kappa) \cdot \nabla v + \sum_{i=1}^{2} \int_{\Omega_i} \kappa_i \nabla u_N(t^j,\kappa) \cdot \nabla v = \int_{\Gamma_{\text{in}}} v, \quad \forall v \in X_N. \tag{33}$$

The RB output is given by $s_N(t^j, \kappa) = \ell(u_N(t^j, \kappa))$ for $0 \le j \le J$. We generate X_N via the Greedy algorithm [21]: we require $\dim(X_N) = 40$ in order to satisfy an error bound tolerance of 0.002 on an RB "training set" — a uniform grid of 900 training parameters in $\mathcal{D}_\kappa$. We note that, for given $\kappa \in \mathcal{D}_\kappa$, in the Online stage — relevant to many queries — it is 109× faster to evaluate the RB ($N = 40$) output *and* output error bound, $\kappa \to s_N(t^j, \kappa), \Delta_N(t^j, \kappa)$, than to directly evaluate the truth, $\kappa \to s_h(t^j, \kappa)$. Recall that $|s_h(t^j, \kappa) - s_N(t^j, \kappa)| \le \Delta_N(t^j, \kappa)$, $0 \le j \le J$, $\kappa \in \mathcal{D}_\kappa$; $\Delta_N(t^j, \kappa)$ is a rigorous error bound.

We now apply our statistcal framework to this proposed parametrized PDE model. We consider two cases below. In the first case we assume that the physical system (PS) is fully modeled by the proposed PDE; we perform validation hypothesis tests to determine which values of the conductivities κ are consistent with (synthetic) experimental data. In the second case we consider "unmodeled physics"; we introduce a physical system (PS′) which does not correspond to the proposed parametrized PDE for any choice of parameters. Note we of course assume that we are unaware of the "unmodeled physics": the validation procedure should

inform us that the physical system PS′ can not be modeled by the proposed PDE.

Fully Modeled Physics

We must first "construct" our Physical System PS. Throughout this section we set $s^{\rm PS}(t^j) \equiv s_h(t^j, \kappa^*)$ for the reference parameter $\kappa^* = (4, 1)$; we then generate synthetic realizations of the form $V^{\rm exp}(t^j, \omega) \equiv s^{\rm PS}(t^j) + \epsilon G(t^j, \omega), 0 \le j \le J$. Here the Gaussian process $G(t, \omega)$ is synthesized from the Karhunen-Loève expansion (5) with correlation length either $\Lambda_\nu = 0$ or $\Lambda_\nu > 0$ (with $C(t^j, t^{j'})$ given by (6) or (7), respectively). We consider the proposed PDE (30) so that the RB output s_N is determined by (33), and our hypothesis tests are based upon $Y^{\hat\kappa}(t^j, \omega) = [s^{\rm PS}(t^j) - s_N(t^j, \hat\kappa)] + \epsilon G(t^j, \omega)$. We consider $I = 0$ — a single Legendre mode — unless otherwise indicated.

We first demonstrate the effect of the choice of N on the possibility region. We set $\epsilon = 0.005$ in (25) and generate $K = 10$ realizations of $V(t_j, \kappa^*, \omega)$, $0 \le j \le J$, for a Gaussian process G which is temporally uncorrelated, $\Lambda_\nu = 0$; we consider $M = 101$ samples in time. We then perform hypothesis tests on a 100×100 uniform grid of parameters in $\mathcal{D}_\kappa$. Figure 2 shows the ($\gamma = 0.95$) possibility region for $N = 10$ and $N = 40$. The reduction in B which results from an increase in the fidelity of the RB model leads to a much sharper characterization of the parametric uncertainty in the problem. We shall henceforth employ only the $N = 40$ RB space.

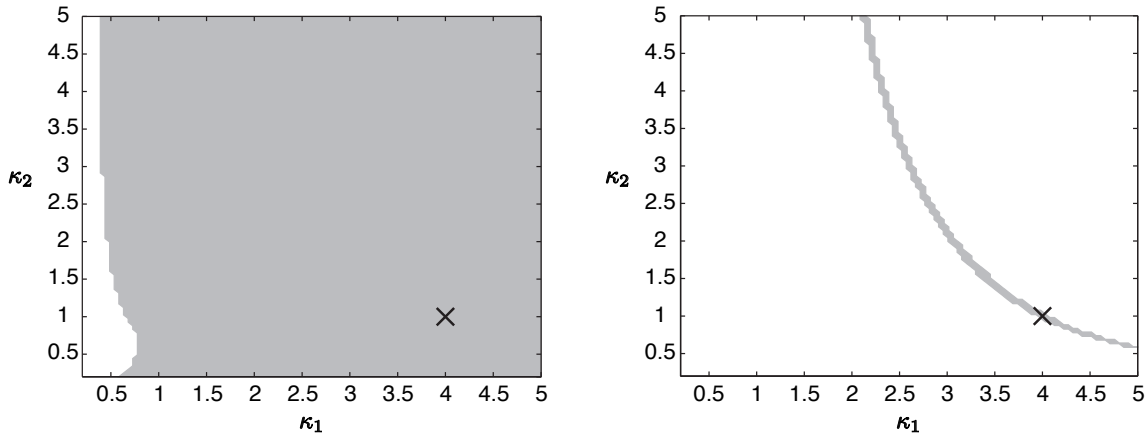

Figure 2: Possibility regions for $\epsilon = 0.005$, $\gamma = 0.95$, $K = 10$, $M = 101$, $I = 0$ and (a) $N = 10$ and (b) $N = 40$; $\kappa^* = (4, 1)$ is indicated in each plot with a green cross.

We next consider the effect of K and M, still with uncorrelated noise. Figure 3a is a repeat of Figure 2a to provide a baseline for $K = 10$ and $M = 101$. In Figure 3b we consider $K = 2$ and $M = 101$: the reduction in K has little effect here because the noise is uncorrelated — the large M compensates for the

small K. Figure 3c presents the possibility region for $K = 2$ and $M = 6$, and in this case the possibility region expands signficantly because we do not have enough data samples. In short, for white noise, large M leads to variance reduction because our calculation of the "moment" $\overline{Z}_K^{[0]}$ is based upon M independent samples.

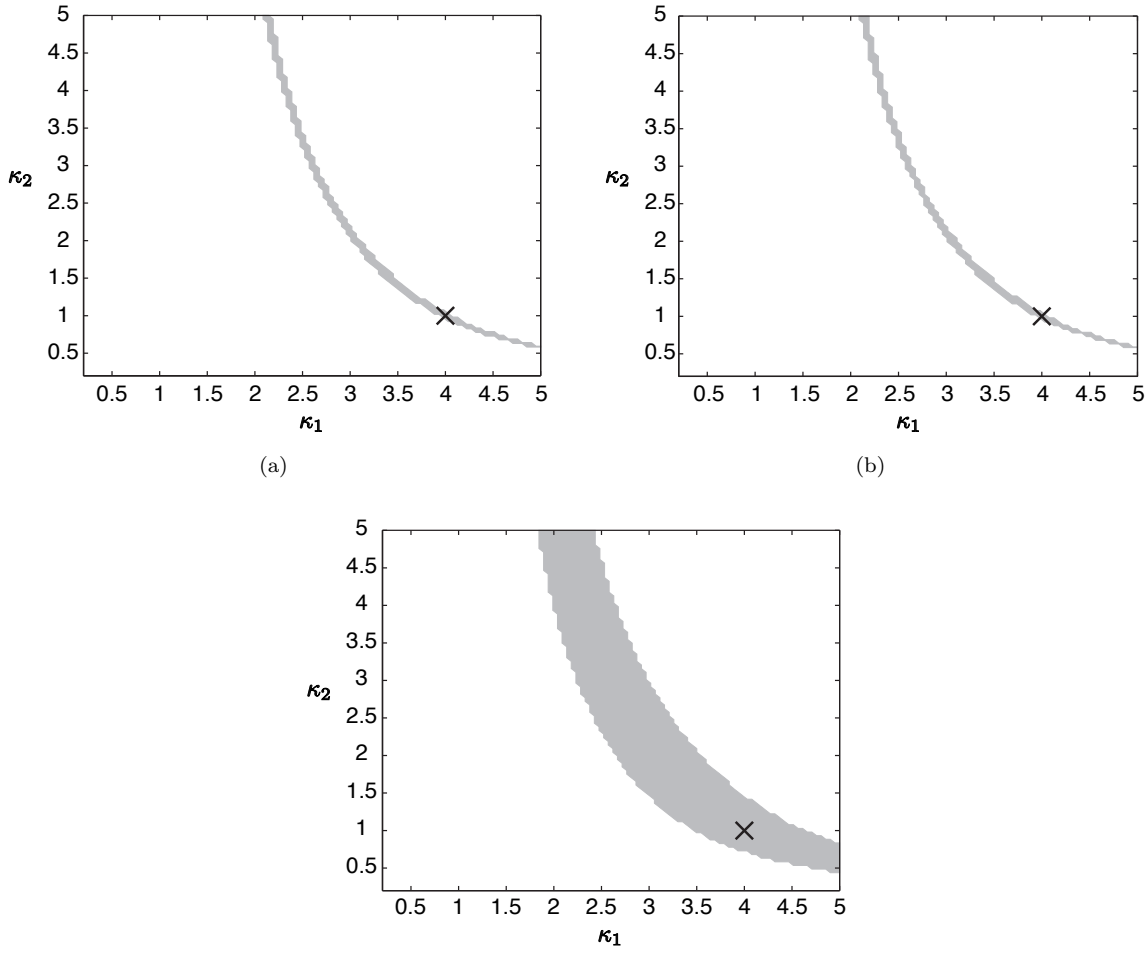

(a) (b)

(c)

Figure 3: Possibility regions for white noise with $\epsilon = 0.005$. We set $\gamma = 0.95$, $I = 0$, $N = 40$ and (a) $K = 10$, $M = 101$; (b) $K = 2$, $M = 101$; (c) $K = 2$, $M = 6$.

We repeat the above study of the effects of K and M but now with *correlated* noise defined according to (7) with $\Lambda_\nu = 1$. We now observe in this correlated noise case — and unlike the white noise case — that for $K = 2$, Figure 4b, the possibility region is significantly larger than for $K = 10$, Figure 4a. Here the M temporal samples within a given "experiment" (or realization) are correlated and therefore, compared to the white noise case of Figure 3, we require more independent realizations (larger K) to reduce the variance of $\overline{Z}_K^{[0]}$.

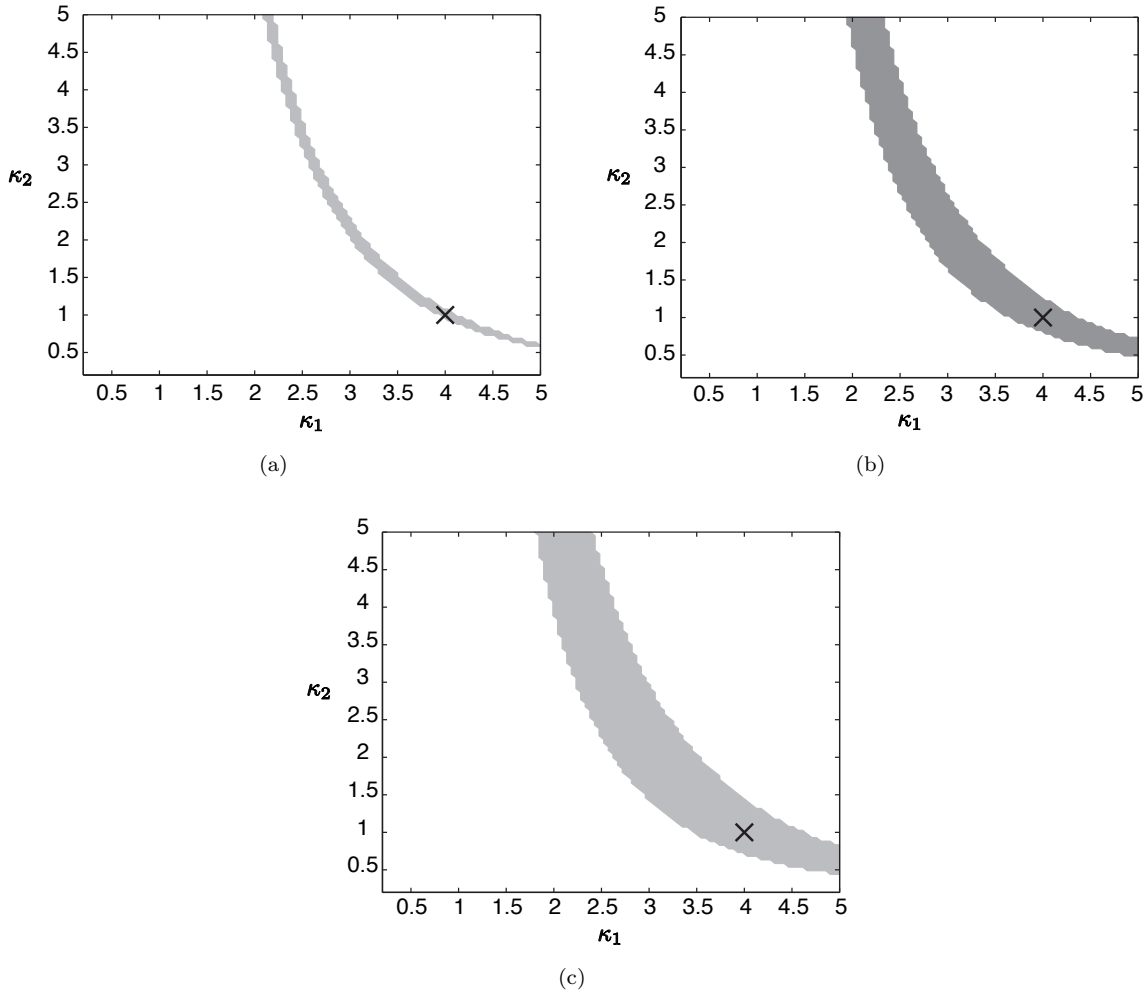

(a)

(b)

(c)

Figure 4: Possibility regions for correlated noise ($\Lambda_t = 1$) with $\epsilon = 0.005$. We set $\gamma = 0.95$, $I = 0$, $N = 40$ and (a) $K = 10$, $M = 101$; (b) $K = 2$, $M = 101$; (c) $K = 2$, $M = 6$.

Finally, we return to the white noise case and demonstrate the effect of changes in the choice of I. We observe from Figures 5a ($I = 0$) and 5b ($I = 4$) that for $I = 4$ we obtain a much sharper characterization of the possibility region — thanks to the higher fidelity characterization of $\delta(t)$. However, Figure 5c shows that for $I = 8$ we again increase the size of the possibility region: the culprit is the denominator of (19); we would need more independent experiments (larger K) in order to accurately characterize the ($I = 8$) 9-variate normally distributed data. We note from our "best result" Figure 5b that there is greater relative uncertainty in (less sensitivity of the output to) κ_2: the first layer enjoys undivided attention for the shorter times during which the thermal penetration depth has not yet reached the second layer. Note that $\epsilon = 0.005$ corresponds to a relative output noise of roughly 1%.

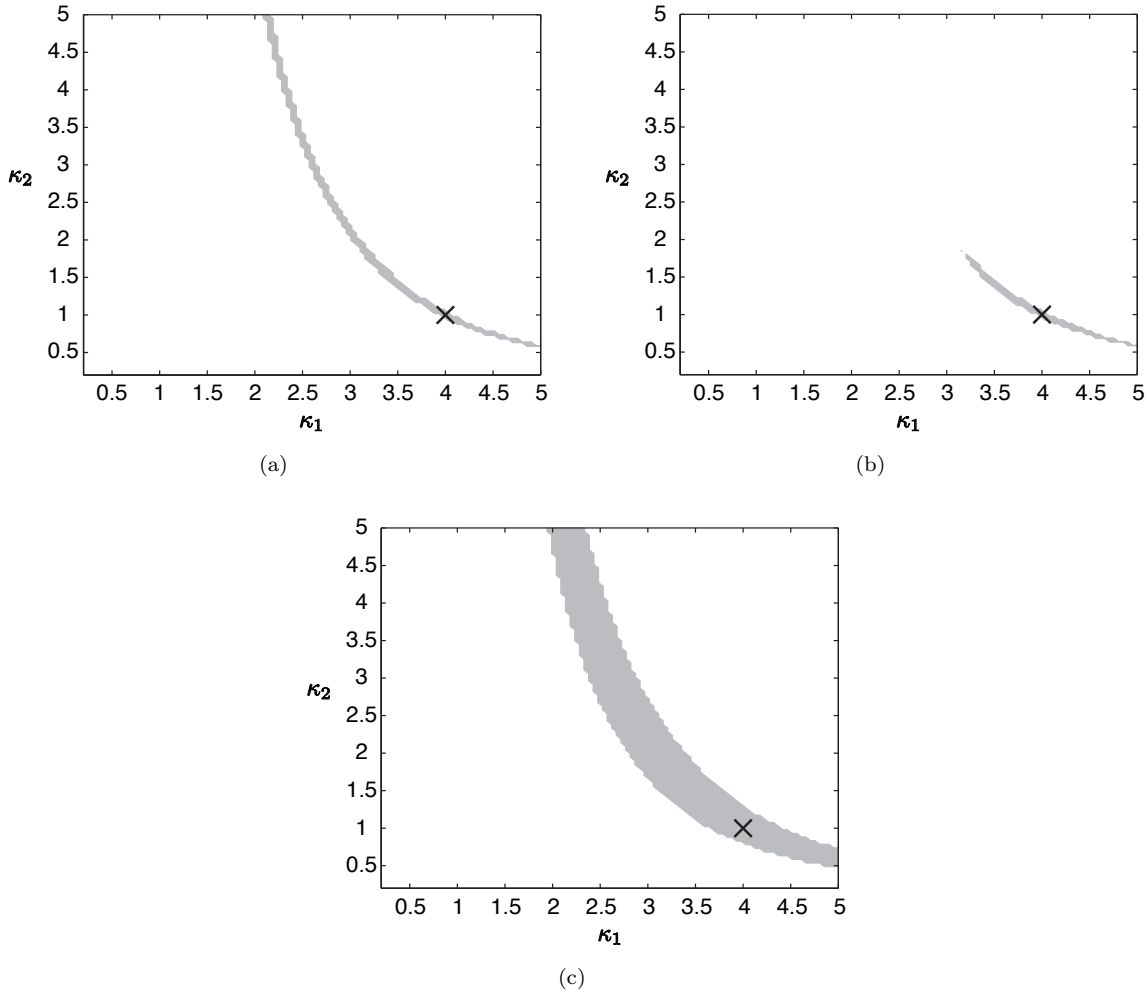

(a) (b)

(c)

Figure 5: Possibility regions for white noise with $\epsilon = 0.005$. We set $\gamma = 0.95$, $K = 10$, $M = 101$, $N = 40$ and (a) $I = 0$; (b) $I = 4$; (c) $I = 8$.

Unmodeled Physics

As already indicated, in our unmodeled physics scenario the proposed parametrized PDE is described, as before, by (30). In order to construct our unmodeled physics and synthetic "experimental" data, however, we shall require a new PDE model which, of course, must include the unmodeled physics — in our case a crack: this new PDE model, described in Figure 6, shall carry a superscript "cracked". In this "cracked" PDE model we introduce a crack of length L_c between Ω_1 and Ω_2. We apply zero flux boundary conditions on the surface of the crack; the weak form is then given by (30) except that the domain Ω is replaced by Ω' (the zero flux condition is natural). Now $s_h^{\text{crack}}(t^j, \kappa^*, L_c)$ shall denote the truth finite element output — defined in the same way as (31) except now the integrals are performed over Ω' — for the "cracked" PDE system with crack length L_c. We now set the output of a "cracked" Physical System, PS$'$, to be

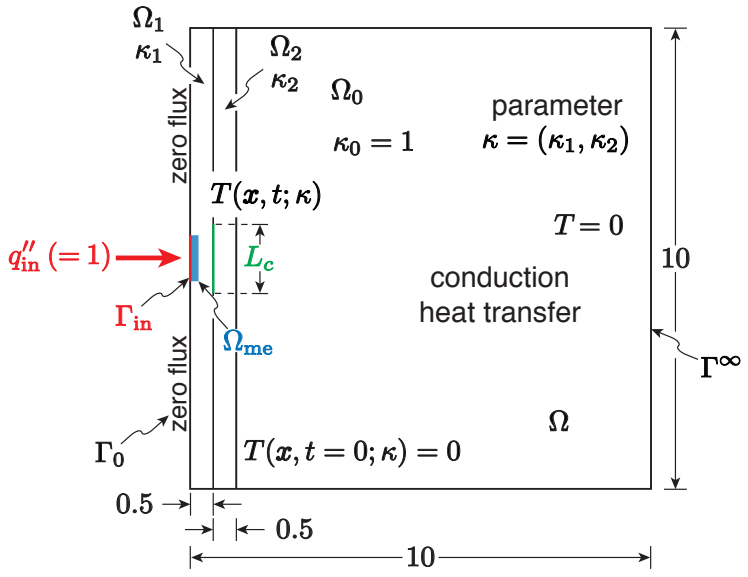

Figure 6: Transient thermal conduction in a domain which contains a crack of length L_c between Ω_1 and Ω_2.

$s^{\text{PS}'}(t^j) \equiv s_h^{\text{crack}}(t^j, \kappa^*, L_c)$; we again set $\kappa^* = (4, 1)$. The experimental realizations are now generated by $V^{\text{exp}}(t^j, \omega) \equiv s^{\text{PS}'}(t^j) + \epsilon G(t^j, \omega), 0 \le j \le J$, where the Gaussian process $G(t^j, \omega)$ is calculated according to (5). We again consider the *proposed* PDE (30) so that the RB output s_N is determined by (33), hence $Y^{\hat{\kappa}}(t^j, \omega) = [s^{\text{PS}'}(t^j) - s_N(t^j, \hat{\kappa})] + \epsilon G(t^j, \omega)$ — we perform model validation with respect to the baseline "idealized" system to test for the presence of pertubations or defects.

Figure 7 shows the possibility regions with cracks of length $L_c = 1$ and $L_c = 1.5$. Figure 7a illustrates that the $L_c = 1$ crack results in a shift in the possibility region so that it no longer contains κ^*; we note that the shift is predominantly toward lower κ_1 as the crack results in a reduction in heat penetration through Ω_1. Figure 7b shows that with $L_c = 1.5$ we obtain an empty possibility region; hence in this case, with a confidence level of $\gamma = 0.95$, we have deduced the presence of unmodeled physics: (32) is not a good model for PS$'$. We note that for sufficiently large N, K, and I it would also be possible to obtain an empty possibility region with the smaller $L_c = 1$ crack, though of course this would increase the computational cost of our procedure.

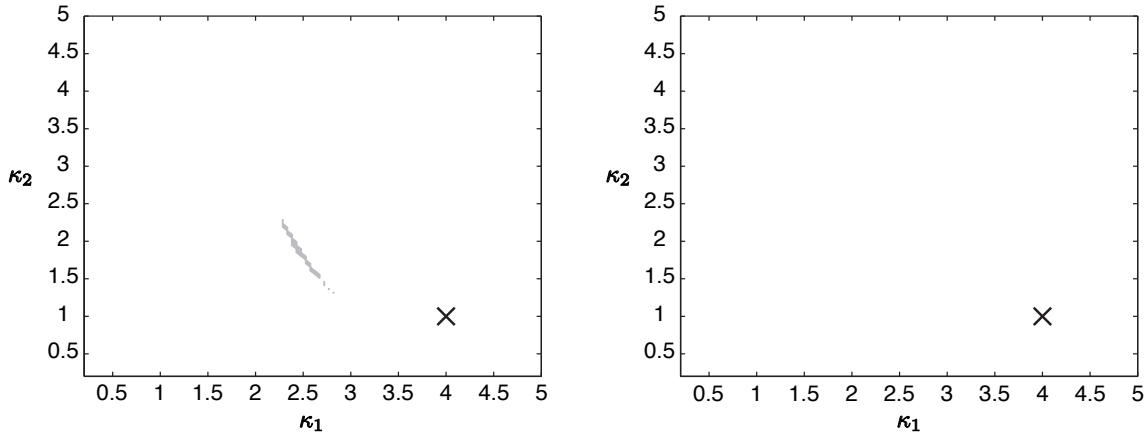

Figure 7: Possibility regions for white noise with $\epsilon = 0.005$. We set $\gamma = 0.95$, $K = 10$, $M = 101$, $N = 40$, $I = 4$ and (a) $L_c = 1$; (b) $L_c = 1.5$.

4.2. Helmholtz Acoustics

Proposed PDE Model

We consider a Helmholtz acoustics impedance model in the domain Ω shown in Figure 8. In the proposed parametrized PDE model, [we choose as parameters the frequency of the (time-harmonic) sound field, k,] and the complex impedance of the wall $Z = Z_r + iZ_i$ ($i = \sqrt{-1}$); we provide a sound source at a speaker diaphragm, and we take as our output the average pressure at the microphone. Similar to the previous example, the proposed parametrized PDE model shall also include both the equation for the field and the prescription of the output.

The relevant domain boundaries are the speaker source (Γ_{in}), the wall of unknown impedance (Γ_{w}), the truncated exterior (Γ^{rad}), and finally the microphone Γ^{mic}. The speaker diaphragm is of length $a = 1.0$, the wall is of length $L = 12$, and the semicircular truncated boundary is of radius $R^{\text{rad}} = 6$ (measured relative to the midpoint of the wall); the distance from the diaphragm to the wall is 4, and Γ^{mic} is halfway between the diaphragm and the wall. Our ν in Section 2 is now k with $\mathcal{D}_k \equiv [1, 2]$, and κ of Section 3 is now (Z_r, Z_i) (impedance, abbreviated as Z) with $(Z_r, Z_i) \in \mathcal{D}_Z \equiv [1, 4] \times [0, 1]$.

We first consider $\mathcal{M}^{\nu,\kappa}$, the weak form of the governing PDE: for $(k, Z) \in \mathcal{D}_k \times \mathcal{D}_Z$, find the acoustic pressure $\phi(k, Z) \in X$ that satisfies

$$\int_\Omega \nabla\phi(k,Z)\cdot\nabla\overline{v} - \int_\Omega k^2\phi(k,Z)\overline{v} + \int_{\Gamma_{\text{w}}} \frac{ik}{Z_r + iZ_i}\phi(k,Z)\overline{v} + \int_{\Gamma^{\text{rad}}} \left(ik + \frac{1}{2R^{\text{rad}}}\right)\phi(k,Z)\overline{v} = -\int_{\Gamma_{\text{in}}} ik\overline{v}, \quad \forall v \in X, \tag{34}$$

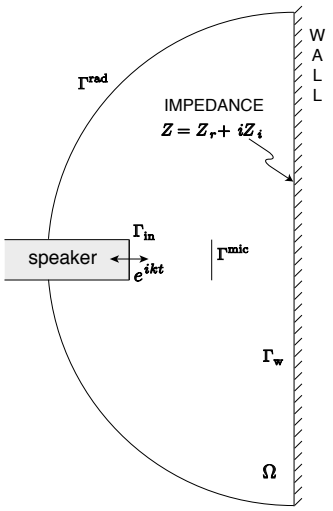

Figure 8: Acoustics impedance problem.

where X is the Hilbert space $X \equiv \{v \in H^1(\Omega)\}$ (over complex-valued fields) and $\overline{v}$ denotes the complex conjugate of v. We note that the third and fourth terms of the left-hand side, and the right-hand side of (34), represent the impedance boundary condition on the wall $\Gamma_{\rm w}$, the first-order radiation boundary condition on the truncated exterior boundary $\Gamma^{\rm rad}$, and the source at the speaker diaphragm $\Gamma_{\rm in}$, respectively. We then consider the output

$$s(k,Z) \equiv \ell(\phi(k,Z)) = \frac{1}{|\Gamma^{\rm mic}|}{\rm Re}\left\{ \int_{\Gamma^{\rm mic}} \phi(k,Z) \right\}, \quad \forall v \in X, \tag{35}$$

which measures the real part (Re) of the averaged pressure over the (microphone) strip $\Gamma^{\rm mic}$.

To obtain $\mathcal{M}_h^{\nu,\kappa}$ we now replace X with a truth finite element space $X_h \subset X$ of dimension $\dim(X_h) = 14126$ degrees of freedom: For $Z \in \mathcal{D}_Z$ and $k \in \mathcal{D}_k$, $\phi_h(k,Z)$ satisfies

$$\int_\Omega \nabla\phi_h(k,Z)\cdot\nabla\overline{v} - \int_\Omega k^2\phi_h(k,Z)\overline{v} + \int_{\Gamma_{\rm w}} \frac{ik}{Z_r + iZ_i}\phi_h(k,Z)\overline{v} + \int_{\Gamma^{\rm rad}} \left(ik + \frac{1}{2R^{\rm rad}}\right)\phi_h(k,Z)\overline{v} = -\int_{\Gamma_{\rm in}} ik\overline{v}, \quad \forall v \in X_h. \tag{36}$$

The truth output is then given by $s_h(k,Z) = \ell(\phi_h(k,Z))$.

We then denote the RB approximation space as X_N and define our RB approximation: For $Z \in \mathcal{D}_Z$ and

$k \in \mathcal{D}_k$, $\phi_N(k,Z)$ satisfies

$$\int_\Omega \nabla\phi_N(k,Z)\cdot\nabla\overline{v} - \int_\Omega k^2\phi_N(k,Z)\overline{v} + \int_{\Gamma_{\rm w}} \frac{ik}{Z_r+iZ_i}\phi_h(k,Z)\overline{v} + \int_{\Gamma^{\rm rad}} \left(ik + \frac{1}{2R^{\rm rad}}\right)\phi_N(k,Z)\overline{v} = -\int_{\Gamma_{\rm in}} ik\overline{v}, \quad \forall v \in X_N; \tag{37}$$

our RB output is then given by $s_N(k,Z) = \ell(\phi_N(k,Z))$. We generate the RB space X_N by the Greedy algorithm: we require $\dim(X_N) \equiv N = 35$ in order to satisfy an error bound tolerance of 10^{-5}. We note that, for given $(k,Z) \in \mathcal{D}_k \times \mathcal{D}_Z$, in the Online stage – relevant to many queries – it is roughly $1000\times$ faster to evaluate the RB ($N = 35$) output and output error bound, $(k,Z) \rightarrow s_N(k,Z)$, $\Delta_N(k,Z)$, than to directly evaluate the truth, $(k,Z) \rightarrow s_h(k,Z)$. Recall that $|s_h(k,Z) - s_N(k,Z)| \leq \Delta_N(k,Z)$, $(k,Z) \in \mathcal{D}_k \times \mathcal{D}_Z$; $\Delta_N(k,Z)$ is a rigorous error bound.

We now apply our statistical framework to this proposed PDE model. We also consider two cases, as in the previous example: a "fully modeled physics" case in which we assume that the physical system (PS) is fully modeled by the proposed PDE model and validation hypothesis tests are performed to determine which value of the impedance is consistent with (synthetic) experiment data; and a "unmodeled physics" case in which a physical system (PS$'$) is introduced which does not correspond to the proposed parametrized PDE for any choice of parameters. Note we of course assume that we are unaware of the "unmodeled physics": the validation procedure should inform us that the physical system (PS$'$) can not be modeled by the proposed PDE. The confidence level for both cases is $\gamma = 0.95$.

In what follows, we shall employ a similar approach as in the previous section except that now each realization corresponds to measurements at M different frequencies (rather than different time levels). The interval $k \in [1,2]$ is divided into $J = 7$ subintervals of equal length $\Delta k = 1/7$ and we set $k^j = 1 + j\Delta k$, $0 \leq j \leq J$; measurements are performed for each frequency k^j, $0 \leq j \leq J$.

Fully Modeled Physics

We first "construct" our Physical System PS. We set $s^{\rm PS}(k^j) \equiv s_h(k^j, Z^*)$ for the reference parameter $Z^* = (Z_r^*, Z_i^*)$; we then generate realizations $V^{\rm exp}(k^j,\omega) = s^{\rm PS}(k^j) + \epsilon G(k^j,\omega)$, $0 \leq j \leq J$, synthesized from the Karhunen-Loève expansion (5) with white noise. We consider the proposed PDE (34) so that the RB output s_N is determined by (37) with $N = 35$, and our hypothesis tests (for any given candidate $\hat{Z}$) are based upon $Y^{\hat{Z}}(k^j,\omega) = [s^{\rm PS}(k^j) - s_N(k^j;\hat{Z})] + \epsilon G(k^j,\omega)$. For all the results below, we consider $\gamma = 0.95$, $K = 8$ realizations, $M = J + 1 = 8$ samples in frequency, and $I = 3$ corresponding to $I + 1 = 4$ Legendre modes. Our possibility region is constructed on a uniform 100×100 grid of parameters in $\mathcal{D}_Z$.

In Figure 9a and 9b we present the possibility regions for $\epsilon = 0.06$ and $\epsilon = 0.01$, respectively. We observe

that the reduction in noise level leads to a much sharper characterization of the parametric uncertainty. For this "best" choice of test options (K, M, I, N) we obtain a rather small possibility region for $\epsilon = 0.01$. Note that $\epsilon = 0.01$ corresponds to a relative output noise of roughly 2%.

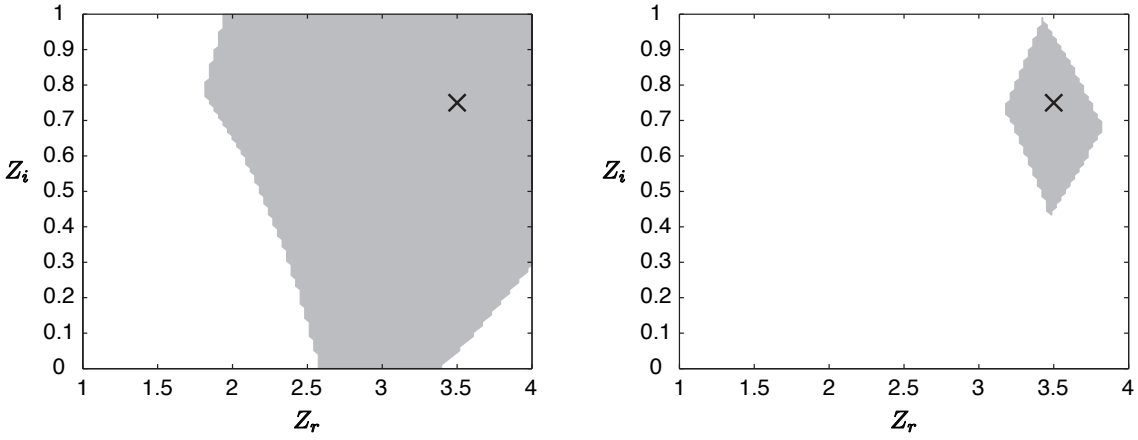

(a) (b)

Figure 9: Possibility regions for (a) $\epsilon = 0.06$ and (b) $\epsilon = 0.01$; $Z^* = (3.5, 0.75)$ is indicated in each plot with a green cross.

Unmodeled Physics

In our unmodeled physics scenario the proposed parametrized PDE is described, as before, by (34). In order to construct our modeled physics and synthetic "experimental" data, we shall define a new PDE which includes the unmodeled physics – in our case, a heterogeneous wall: this new PDE model, which corresponds to the model shown in Figure 10, shall carry a superscript "2-part wall." In particular, we now assume that the wall consists of two parts: the upper part, $\Gamma_{\text{w,top}}$ with real infinite impedance (hard wall); and the lower part, $\Gamma_{\text{w,bot}}$, with impedance $(Z_r, Z_i) \in \mathcal{D}_Z$. The "2-part wall" PDE weak form is thus the same as in (34), except now the integral over Γ_{w} is taken only over the lower half of the wall, $\Gamma_{\text{w,bot}}$. We now set $s^{\text{PS}'}(k^j) \equiv s_h^{\text{2-part wall}}(k^j, Z^*)$, where $s_h^{\text{2-part wall}}(k^j, Z^*)$ denotes the "2-part wall" truth finite element output for $Z^* = (3.5, 0.75)$. Finally, we take $V^{\text{exp}}(k^j, \omega) \equiv s^{\text{PS}'}(k^j) + \epsilon G(k^j, \omega)$, $0 \le j \le J$, where the Gaussian process $G(k, \omega)$ is calculated according to (5). We now pursue our hypothesis tests with $Y^{\hat{Z}}(k^j, \omega) = [s^{\text{PS}'}(k^j) - s_N(k^j; \hat{Z})] + \epsilon G(k^j, \omega)$ for s_N evaluated from (37) for $N = 35$: we perform validation with respect to the baseline homogeneous wall to detect the presence of heterogeneities.

Figure 11 shows the possibility regions with noise levels $\epsilon = 0.06$ and $\epsilon = 0.01$. Figure 11 illustrates that for large noise the difference between V^{exp} and s_N results in a shift in the possibility region such that it no longer contains Z^*; however, for small noise, we obtain an empty possibility region. In the former case

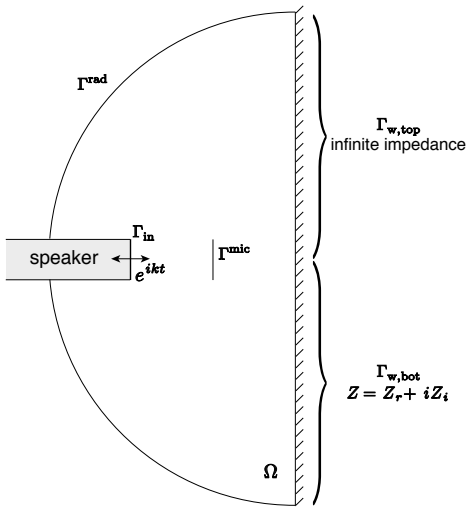

Figure 10: The "2-part wall" acoustics impedance problem.

the test (incorrectly) reconciles the data through an artificially large impedance. In the latter case we have deduced the presence of unmodeled physics – a wall with non-uniform impedance.

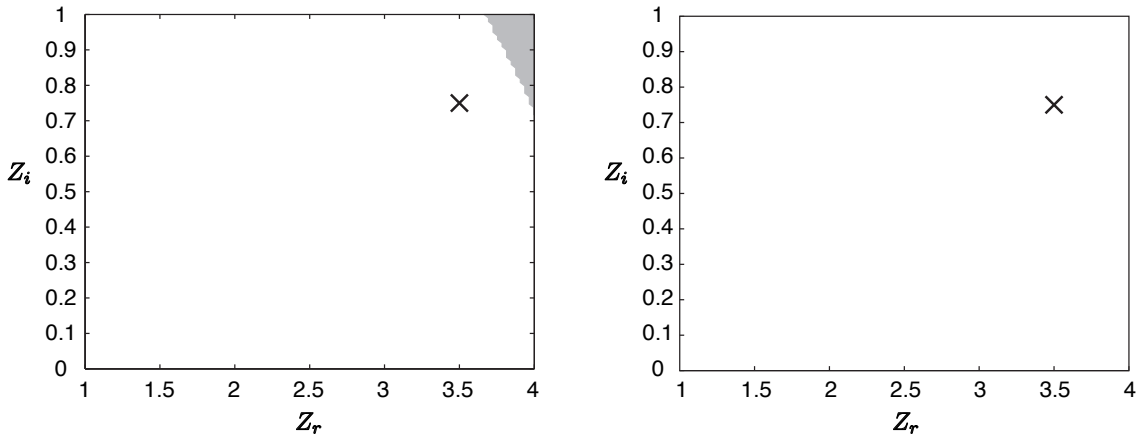

(a) (b)

Figure 11: Possibility regions for (a) $\epsilon = 0.06$ and (b) $\epsilon = 0.01$.

5. Acknowledgments

We would like to thank Professor Youssef Marzouk of MIT for many fruitful discussions. This work was supported by OSD/AFOSR/MURI Grant FA9550-09-1-0613 and the MIT-Singapore International Design Center.

References

[1] H. Banks, K. Kunisch, Estimation Techniques for Distributed Parameter Systems, Birkhäuser, 1989.

[2] A. Tarantola, Inverse Problem Theory and Model Parameter Estimation, 2005.

[3] T. Trucano, L. Swilera, T. Igusab, W. Oberkampf, M. Pilch, Calibration, validation, and sensitivity analysis: Whats what, Reliab. Engrg. Syst. Safe. 91 (2006) 1331–1357.

[4] M. Kennedy, A. O'Hagan, Bayesian calibration of computer models, J. Roy. Stat. Soc. Ser. B 63 (2001) 425–464.

[5] Y. Xiong, C. W., K.-L. Tsui, D. Apley, A better understanding of model updating strategies in validating engineering models, Comput. Methods Appl. Mech. Engrg. (2009) 1327–1337.

[6] P. B. Stark, L. Tenorio, A primer of frequentist and bayesian inference in inverse problems, in: Biegler, Biros, Ghattas, Heinkenschloss, Keyes, Mallick, Tenorio, van Bloemen Waanders, Willcox (Eds.), Computational Methods for Large Scale Inverse Problems and Uncertainty Quantification, John Wiley & Sons, UK, 2009.

[7] J. McFarland, S. Mahadevan, Multivariate significance testing and model calibration under uncertainty, Comput. Methods Appl. Mech. Engrg. 197 (29–32) (2008) 2467–2479.

[8] O. Balci, R. Sargent, Validation of simulation models via simultaneous confidence intervals, Am. J. Math. Manage. Sci. 4 (1984) 375–406.

[9] A. Mood, F. Graybill, D. Boes, Introduction to the Theory of Statistics, third edition Edition, McGraw-Hill, 1974.

[10] V. Chew, Confidence, prediction, and tolerance regions for the multivariate normal distribution, Journal of the American Statistical Association 61 (315) (1966) 605–617.

[11] K. Dowding, M. Pilch, R. Hills, Formulation of the thermal problem, Comput. Methods Appl. Mech. Engrg. 197 (29–32) (2008) 2385–2389.

[12] A. Quarteroni, A. Valli, Numerical Approximation of Partial Differential Equations, Springer, 2008.

[13] N. Draper, H. Smith, Applied Regression Analysis, 3rd Edition, Wiley, 1998.

[14] P. Athanasios, S. U. Pillai, Probability, Random Variables and Stochastic Processes, McGraw-Hill, 2001.

[15] P. Ciarlet, The Finite Element Method for Elliptic Problems, North-Holland, 1978.

[16] T. A. Porsching, M. Y. L. Lee, The reduced-basis method for initial value problems, SIAM Journal of Numerical Analysis 24 (1987) 1277–1287.

[17] J. P. Fink, W. C. Rheinboldt, On the error behavior of the reduced basis technique for nonlinear finite element approximations, Z. Angew. Math. Mech. 63 (1) (1983) 21–28.

[18] B. O. Almroth, P. Stern, F. A. Brogan, Automatic choice of global shape functions in structural analysis, AIAA Journal 16 (1978) 525–528.

[19] A. K. Noor, Recent advances in reduction methods for nonlinear problems, Comput. Struct. 13 (1981) 31–44.

[20] G. Rozza, D. B. P. Huynh, A. T. Patera, Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations, Archives Computational Methods in Engineering 15 (3) (2008) 229–275.

[21] K. Veroy, C. Prud'homme, D. V. Rovas, A. T. Patera, *A Posteriori* error bounds for reduced-basis approximation of parametrized noncoercive and nonlinear elliptic partial differential equations, in: Proceedings of the 16th AIAA Computational Fluid Dynamics Conference, 2003, paper 2003-3847.

[22] J. Wang, N. Zabaras, Using Bayesian statistics in the estimation of heat source in radiation, International Journal of Heat and Mass Transfer 48 (2005) 15–29.

[23] G. R. Liu, J. H. Lee, A. T. Patera, Z. L. Yang, K. Y. Lam, Inverse identification of thermal parameters using reduced-basis method, Comput. Methods Appl. Mech. Engrg. (27–29) (2005) 3090–3107.

[24] N. C. Nguyen, G. Rozza, D. B. P. Huynh, A. T. Patera, Reduced basis approximation and a posteriori error estimation for parametrized parabolic PDEs; Application to real-time Bayesian parameter estimation, in: L. Biegler, G. Biros, O. Ghattas, M. Heinkenschloss, D. Keyes, B. Mallick, Y. Marzouk, L. Tenorio, B. van Bloemen Waanders, K. Willcox (Eds.), Large-Scale Inverse Problems and Quantification of Uncertainty, John Wiley & Sons, Ltd.

[25] M. Loève, Probability theory, 4th Edition, Vol. 2, Springer-Verlag, 1978.

[26] E. H. Isaaks, R. M. Srivastava, An Introduction to Applied Geostatistics, Oxford University Press, 1989.

[27] M. Abramowitz, I. Stegun (Eds.), Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables, Dover Publications, 1972.

[28] M. A. Grepl, A. T. Patera, A posteriori error bounds for reduced-basis approximations of parametrized parabolic partial differential equations, Mathematical Modelling and Numerical Analysis 39 (1) (2005) 157–181. doi:10.1051/m2an:2005006.

[29] B. Haasdonk, M. Ohlberger, Reduced basis method for finite volume approximations of parametrized evolution equations 42 (2) (2008) 277–302.

[30] D. J. Knezevic, A. T. Patera, A certified reduced basis method for the Fokker-Planck equation of dilute polymeric fluids: FENE dumbbells in extensional flow, SIAM Journal on Scientific Computing 32 (2) (2010) 793–817.

[31] J. Nocedal, S. J. Wright, Numerical Optimization, 2nd Edition, Springer, 2006.

[32] B. S. Kirk, J. W. Peterson, R. M. Stogner, G. F. Carey, libMesh: A C++ library for parallel adaptive mesh refinement/coarsening simulations, Engineering with Computers 23 (3–4) (2006) 237–254.

[33] D. J. Knezevic, J. W. Peterson, A high-performance implementation of the certified reduced basis method, Accepted to Comput. Methods Appl. Mech. Engrg.