

Tikhonov regularization as a nonparametric method for uncertainty quantification in aggregate data problems

Elena Villalón^a, Qian Yang^b, Carlos A. Sing Long^{a,c,*}

^a Institute for Mathematical and Computational Engineering, Pontificia Universidad Católica de Chile, Av. Vicuña Mackenna 4860, Santiago, 7820436, Chile

^b School of Computing, University of Connecticut, 371 Fairfield Way, Unit 4155, Storrs, 06269, CT, USA

^c Institute for Biological and Medical Engineering, Pontificia Universidad Católica de Chile, Av. Vicuña Mackenna 4860, Santiago, 7820436, Chile



ARTICLE INFO

Keywords:

Uncertainty quantification
Dynamical systems
Convex optimization
Aggregate data problems
Regularization
Nonparametric methods

ABSTRACT

Dynamical systems are widely used in mathematical models in engineering and in the applied sciences. However, the parameters of these systems are usually unknown and they must be estimated from measured data. Performing parameter estimation and quantifying the uncertainty in these estimates becomes critical to realize the predictive power of dynamical systems. In some applications, only measurements about the states of an ensemble of systems are available, where each one evolves according to the same model but for different parameter values, leading to *aggregate data* problems. In this case, an approach proposed in the literature is to estimate a density that is consistent with some estimates of the expected value of some quantities of interest. To solve the problem in practice, the density is discretized and, to solve the resulting ill-posed problem, Tikhonov regularization is used. In this work, we propose a Bayesian model that shows this approach can be interpreted as a *maximum a posteriori* estimate for the density. We show that the infinite-dimensional problem defining the MAP can be reformulated as a finite-dimensional problem that does not require discretizing the density. In several cases of interest, this problem is convex, unconstrained, and the objective function is smooth. Thus, it can be solved using algorithms with optimal convergence rates. The trade-off is that the objective is defined by an integral. However, our results characterize the regularity of the integrand, allowing the use of tailored numerical schemes to approximate it. Furthermore, our theoretical results characterize the form of the optimal density, whereas our numerical results illustrate the performance of our method and confirm our theoretical findings.

1. Introduction

Dynamical systems are a cornerstone of mathematical modeling in engineering and in the applied sciences. Models based on dynamical systems are broadly used due to their predictive power [1,2] and due to the availability of numerical methods that accurately simulate them [3–5]. However, in practical applications these models depend on unknown parameters that must be estimated from some observed trajectories of the system. Thus, to realize the predictive power of dynamical systems in practice, it

* Corresponding author.

E-mail address: casinglo@uc.cl (C.A. Sing Long).

is not only important to develop methods to estimate these parameters from experimental measurements, but also to quantify the uncertainty in these estimates, and their impact on the predictions obtained from numerical simulations [6–9].

Typically, several *quantities of interest* (QoIs) are measured, possibly with errors, from observed trajectories. These measurements are then used to estimate the unknown parameters and to quantify the uncertainty on these estimates, i.e., the objective is to perform *model calibration* and/or *inverse uncertainty quantification* [8, Chapter 1]. If we have measurements of the QoIs from several trajectories generated by the same system, i.e., a system for which there is no uncertainty in the parameters, we obtain a *longitudinal data problem* [10]. In this case, the likelihood can usually be characterized and Bayesian methods are well-suited to simultaneously estimate the unknown parameters, e.g., using the posterior expectation or the maximum *a posteriori* (MAP), and to quantify the uncertainty in these estimates (see, e.g., [8, Chapter 8]). However, in some applications one can only measure the QoIs from trajectories generated by an ensemble of systems that evolve according to the same model, but may do so according to different values of the unknown parameters. As a concrete example, estimating the unknown parameters in the CAR T-Cell cancer model proposed by Schacht et al. [11] requires measuring the amount of T-cells in mice. As this measurement cannot be performed in the same mouse at different times, a cohort is used, and one measurement is performed for each mouse. Thus, we have observations about systems that evolve according to the same model, but for different values of the unknown parameters. In this case, the parameters are assumed to be random variables with an unknown distribution, and we obtain what is called an *aggregate data problem* (see, e.g., [12] or [13, Chapter 5]). Thus, the natural object to estimate is the unknown probability distribution of the parameters, and to quantify the uncertainty in this estimate.

Banks et al. [14] propose to estimate the probability density by enforcing consistency with the expected value of a QoI estimated from several measurements made about an ensemble of systems. Since the object to be estimated is a density, and thus an infinite-dimensional object, this is a *nonparametric method*. To compute this density in practice, they propose to discretize it, e.g., using a mesh, splines, or other approximations, to then solve a least-squares problem enforcing consistency with the measured data (see, e.g., [13, Chapter 5] or [15, Chapter 14]). The discretization leads to a parametric method which may lead to determined or underdetermined problems depending on the number of constraints imposed on the density and the number of degrees of freedom used to discretize it. The authors in [16] extended this approach by enforcing additional constraints on the unknown density, such as constraining the value of its higher-order moments, or the dispersion of the state of the system around a given point at specific times. To compute a discretized density satisfying these constraints efficiently, the authors leverage convex programming and the Koopman operator. This relies on the fact that the method proposed by Banks et al. leads to a convex least-squares problem, and on the fact that a single probability density at the initial time that can be pushed forward in time by the system dynamics using the Koopman operator. In fact, when the constraints are linear, they become constraints on the expected value of the Koopman operator applied to a QoI [17]. To regularize the solution to a possibly underdetermined system, which occurs when the degrees of freedom of the discretization exceed the number of constraints, the authors propose penalizing the values of the density at its discretization points by its ℓ_2 -norm squared, i.e., by applying Tikhonov regularization [18]. The resulting problem can be interpreted as *Tikhonov regularization for moment constraints* on the unknown discretized density.

Although this approach has appealing properties, some open questions remain about its practical implementation. On one hand, solving the problem by discretizing the density can implicitly restrict its support, and, when the discretization is over a mesh as suggested in [17], it can preclude the use of the method in several practical applications. There is a growing number of engineering problems where the input to the model and its parameters are high-dimensional. For instance, the quantification of uncertainties in Mars' atmospheric entry, descent, and landing processes due to the uncertainty present in the atmosphere density, the initial states, and other model parameters [19], wind-turbine simulators that take time series of simulated wind speed as the input to predict mechanical loads on its sub-structures and fatigue [20], and the study of variations on fabrication processes on nano-scale chip design, such as the surface roughness of interconnects and the random doping effects in transistors, that influence the performance of chips [21].

On the other hand, Tikhonov regularization has several computational benefits, particularly its use of a smooth and strongly convex regularizer, and it is widely used to solve ill-posed or ill-conditioned inverse problems. However, its use as a regularization strategy to estimate a probability density lacks an intuitive statistical interpretation, particularly in contrast to other regularizers, such as the KL divergence [22]. This limits the appeal of the approach, and makes the interpretation of its results challenging. Consequently, both developing a theoretical framework that allows us to interpret this method as a classical estimation procedure, and developing implementations of this method that are computationally efficient, along with suitable theoretical guarantees, would foster its adoption for uncertainty quantification in aggregate data problems involving high-dimensional dynamical systems.

1.1. Contributions

To introduce the method, consider the case of a dynamical system in \mathbb{R}^d given by

$$\begin{cases} \dot{\mathbf{x}}(t) = \mathbf{f}(t, \mathbf{x}(t)) \\ \mathbf{x}(0) = \mathbf{x}_0 \end{cases}$$

where there is uncertainty in the initial condition. For simplicity, suppose that we have an estimate $\bar{\mathbf{x}}_{t_0}$ of the expected state at a time t_0 . Then, to estimate the uncertainty on the initial condition from observations about the state of the system, we may look for a density ρ that satisfies the constraint

$$\int \mathbf{x}(t_0, \mathbf{x}_0) \rho(\mathbf{x}_0) d\mathbf{x}_0 \approx \bar{\mathbf{x}}_{t_0},$$

where $\mathbf{x} = \mathbf{x}(t, \mathbf{x}_0)$ is the state at time t for the initial condition \mathbf{x}_0 . There is an infinite number of densities satisfying this constraint exactly. The method proposed in [16] involves solving the convex problem

$$\begin{aligned} \text{minimize}_{\rho} \quad & \lambda \int \rho(\mathbf{x}_0)^2 d\mathbf{x}_0 + \frac{1}{2} \left\| \int \mathbf{x}(t_0, \mathbf{x}_0) \rho(\mathbf{x}_0) d\mathbf{x}_0 - \bar{\mathbf{x}}_{t_0} \right\|_2^2 \\ \text{subject to} \quad & \rho \geq 0, \quad \int \rho(\mathbf{x}_0) d\mathbf{x}_0 = 1, \end{aligned} \quad (1)$$

for some regularization parameter $\lambda > 0$ and over a suitable functional space. This problem can be interpreted as Tikhonov regularization on moment constraints. A direct approach to solve this problem numerically in low-dimensions is to discretize the integrals using $N := n^d$ equispaced nodes $\{\mathbf{x}_{0,i}\}$. Then, by defining

$$\rho = \begin{bmatrix} \rho(\mathbf{x}_{0,1}) \\ \vdots \\ \rho(\mathbf{x}_{0,N}) \end{bmatrix} \quad \text{and} \quad \mathbf{a}_i = \begin{bmatrix} \mathbf{x}_i(t_0, \mathbf{x}_{0,1}) \\ \vdots \\ \mathbf{x}_i(t_0, \mathbf{x}_{0,N}) \end{bmatrix}$$

we may solve

$$\begin{aligned} \text{minimize}_{\rho} \quad & \frac{\lambda \Delta x^d}{2} \|\rho\|_2^2 + \frac{1}{2} \sum_{i=1}^d (\Delta x^d \langle \mathbf{a}_i, \rho \rangle_2 - \bar{x}_{t,i})^2 \\ \text{subject to} \quad & \rho \geq 0, \quad \Delta x^d \langle \mathbf{1}, \rho \rangle_2 = 1 \end{aligned}$$

where $\mathbf{1}$ is the vector with all its entries equal to one and Δx is the discretization step. On one hand, this is a convex optimization problem with a number of variables that is exponential in the number of dimensions, precluding the use of this method for high-dimensional problems. On the other hand, it is not intuitive why (1) is an efficient estimation procedure for the density, nor what is the class of possible densities that can be selected by this procedure.

In this work, we introduce a Bayesian model that provides a theoretical framework to interpret (1) as a MAP estimate for the density. This not only relates Tikhonov regularization to nonparametric estimation, in the sense that the parameter lies on an infinite-dimensional space [10], but possibly with Bayesian nonparametric methods [23]. Furthermore, we provide a method to solve (1) without needing to discretize the density. By leveraging convex duality, we prove that, even in general cases, the problem (1) can be reduced to a smooth, finite-dimensional convex optimization problem in $d + 1$ variables which involves the computation of a d -dimensional integral. Furthermore, the dual problem allows us to provide a closed-form expression for the optimal density, a result that has both computational and statistical implications.

We consider the following as the main contributions of our work.

- (i) **Tikhonov regularization as a MAP estimate:** We propose a Bayesian model that provides a rigorous justification for using Tikhonov regularization on moment constraints to perform input uncertainty quantification, and allows us to interpret the computed density as a MAP estimate for this model.
- (ii) **Connections to regularized regression:** When we have incomplete linear measurements about some QoIs, we show that solving (1) is equivalent to solving a regularized regression problem for the expected values of the QoIs themselves.
- (iii) **Geometry and support:** We explicitly study the effect of restricting the support of ρ to a convex set Ω , showing through theoretical results and numerical experiments that it has a significant impact on the estimated density.
- (iv) **Closed-form expressions:** We provide closed-form expressions for the optimal probability density, showing the explicit dependence of the probability density both on the QoIs and on the dynamics of the system.
- (v) **High-dimensional problems:** We show that by solving a finite-dimensional convex problem with a number of variables that is proportional to the number of QoIs, we may recover the optimal solution to (1), enabling the efficient implementation of the method for high-dimensional problems.

1.2. Organization

The paper is organized as follows. In Section 2 we develop a Bayesian model that allows us to interpret Tikhonov regularization as a MAP estimate for the density. In Section 3 we show that the infinite-dimensional convex problem defining the MAP leads to an equivalent finite-dimensional problem. In Section 4 we explore the connections between our method and a regularized regression problem for the expected values of the QoIs. In Section 5 we discuss efficient methods to find an optimal solution to the optimization problem we introduce. In Section 6 we show study the performance of the method in terms of the kind of densities it selects, and the effect of the support on them. We conclude in Section 7 with a discussion of our results.

2. Uncertainty quantification by Tikhonov regularization

2.1. The model

In this work, we consider *random differential equations* [8, Chapter 3] in \mathbb{R}^d of the form

$$\begin{cases} \dot{x}(t) = f(t, x(t)) \\ x(0) = x_0 \end{cases} \quad \text{where } x_0 \sim \Pi_{x_0} \quad (2)$$

for some unknown probability measure Π_{x_0} on \mathbb{R}^d . Typically, it will have a density ρ_{x_0} with respect to a reference measure μ such as the Lebesgue measure. We assume that this equation accurately models the physical phenomenon of interest, that is, that any *model uncertainty* [8, Chapter 2] is negligible. In practical applications, we only have available sample trajectories $x = x(t, x_0)$ of the system for a few realizations of x_0 and for each one of these realizations we typically measure a small number of fixed *quantities of interest* (QoIs) u_1, \dots, u_q . Since the fluctuations in the values of these QoIs are only due to the fluctuations in the initial condition, we represent these QoIs as a function

$$u = U(x_0).$$

We shall assume that the QoIs can be computed to high accuracy and precision, so that there are no significant *numerical uncertainties* [8, Chapter 2] associated to this model. The measurement of the QoIs is typically corrupted by additive noise, which is statistically independent of x_0 . Therefore, instead of u we observe

$$y = U(x_0) + \varepsilon \quad (3)$$

for some centered random vector ε . Thus, a sample $x_0^{(1)}, \dots, x_0^{(n)}$ of unknown initial conditions leads to a sample of QoIs $u^{(1)}, \dots, u^{(n)}$ and measurements $y^{(1)}, \dots, y^{(n)}$ from which we want to quantify the uncertainty on the initial condition, that is, to quantify the *input uncertainty* [8, Chapter 3] in (2). This is sometimes called an *aggregate data inverse problem* [10,12]. In this case, it is not sufficient to provide a point estimate for the initial condition, as it is a random variable. Hence, it is necessary to find a suitable probability distribution that approximates Π_{x_0} and thus the uncertainty on x_0 .

Instead of formulating a Bayesian model directly from (3) we proceed as follows. We assume throughout that $\varepsilon \sim \mathcal{N}(0, \sigma^2 I_q)$. Then

$$y | u \sim \mathcal{N}(u, \sigma^2 I_q).$$

To model the distribution of u we decompose its behavior in terms of an average value \bar{u} and fluctuations around this average that do not depend on exogenous parameters, that is, we assume that

$$u | \bar{u} \sim \Pi_{u|\bar{u}}$$

for some probability measure $\Pi_{u|\bar{u}}$. To model the behavior of the average, let $\{\Pi_{x_0|\theta} : \theta \in \Theta\}$ be a parametric model for the distribution of x_0 with parameter space $\Theta \subset \mathbb{R}^s$ and let

$$\bar{u} | \theta = \mathbf{E}_{x_0 \sim \Pi_{x_0|\theta}}[U(x_0)].$$

Finally, we let Π_θ be a prior for θ . This leads to the Bayesian model

$$\begin{cases} y | u \sim \mathcal{N}(u, \sigma^2 I_q), \\ u | \bar{u} \sim \Pi_{u|\bar{u}}, \\ \bar{u} | \theta = \mathbf{E}_{x_0 \sim \Pi_{x_0|\theta}}[U(x_0)], \\ \theta \sim \Pi_\theta. \end{cases} \quad (4)$$

If $\Pi_{\theta|y}$ is the posterior obtained from this model, then the posterior for x_0 is

$$\mathbf{P}[x_0 \in A] = \int_{\Theta} \Pi_{x_0|\theta}(A) d\Pi_{\theta|y}(\theta)$$

The dimension s of the parameter space Θ constrains the flexibility of the model we use for \bar{u} which, in turn, constraints the flexibility of the posterior for x_0 . For this reason, we will introduce a parameter space Θ and a suitable prior measure Π_θ on Θ that allows us to control the flexibility of the model, focusing on the asymptotic regime for $s \gg 1$. Furthermore, instead of aiming to find, or sample from, the posterior distribution, we will focus on characterizing the *maximum a posteriori* (MAP) for the parameter θ . However, before proceeding, we discuss some extensions of this model and its connections with previous work.

2.1.1. Partial measurements of QoIs

In some applications, we cannot measure the full set of QoIs for the sample trajectories. The results we will present extend naturally to the case in which we observe *incomplete linear measurements* of the QoIs. In other words, we assume that there exists a $m \times q$ matrix \mathbf{A} such that we are only able to measure $\mathbf{A}\mathbf{u}$. This leads to the *partial measurement model*

$$\mathbf{y} = \mathbf{A}\mathbf{U}(\mathbf{x}_0) + \boldsymbol{\varepsilon}.$$

In this case, for $\mathbf{u}' = \mathbf{A}\mathbf{u}$ the model (4) simply becomes

$$\begin{cases} \mathbf{y} | \mathbf{u}' \sim \mathcal{N}(\mathbf{u}', \sigma^2 \mathbf{I}_m) \\ \mathbf{u}' | \bar{\mathbf{u}}' \sim \Pi_{\mathbf{u}'} | \bar{\mathbf{u}}' \\ \bar{\mathbf{u}}' | \boldsymbol{\theta} = \mathbf{A} \mathbf{E}_{\mathbf{x}_0 \sim \Pi_{\mathbf{x}_0} | \boldsymbol{\theta}} [\mathbf{U}(\mathbf{x}_0)] \\ \boldsymbol{\theta} \sim \Pi_{\boldsymbol{\theta}} \end{cases}$$

and the same results that hold for (4) will hold for this model. It is straightforward to see this is equivalent to (3) if we define $\mathbf{U}' = \mathbf{A}\mathbf{U}$. However, in Section 4 we discuss the advantages of treating the measurement matrix \mathbf{A} separately from the QoIs.

2.1.2. Connection with the Koopman operator

A special case of QoIs has the form

$$\mathbf{U}(\mathbf{x}_0) = \begin{bmatrix} U_1(\mathbf{x}(t_1, \mathbf{x}_0)) \\ \vdots \\ U_q(\mathbf{x}(t_n, \mathbf{x}_0)) \end{bmatrix}$$

for some known instants t_1, \dots, t_n and functions U_1, \dots, U_q . In this case, we may represent the QoIs in terms of the *Koopman operator* [24,25]. If we denote as $\Phi^t(\mathbf{x}_0) := \mathbf{x}(t, \mathbf{x}_0)$ the flow of the system [2, Chapter 1] then the Koopman operator K_t maps $f : \mathbb{R}^d \rightarrow \mathbb{R}$ to

$$K_t f(\mathbf{x}_0) := f(\Phi^t(\mathbf{x}_0)).$$

This allows us to write

$$\mathbf{U}(\mathbf{x}_0) = \begin{bmatrix} K_{t_1} U_1(\mathbf{x}_0) \\ \vdots \\ K_{t_q} U_q(\mathbf{x}_0) \end{bmatrix}.$$

The expected value of \mathbf{U} can also be represented in terms of the pushforward of $\Pi_{\mathbf{x}_0}$ under Φ^t at each t_1, \dots, t_q . In fact,

$$\int K_t f(\mathbf{x}_0) d\Pi_{\mathbf{x}_0}(\mathbf{x}_0) = \int f(\Phi^t(\mathbf{x}_0)) d\Pi_{\mathbf{x}_0}(\mathbf{x}_0) = \int f(\mathbf{x}_t) d\Phi_*^t \Pi_{\mathbf{x}_0}(\mathbf{x}_t)$$

or, equivalently, $\mathbf{x}(t, \mathbf{x}_0) \sim \Phi_*^t \Pi_{\mathbf{x}_0}$ when $\mathbf{x}_0 \sim \Pi_{\mathbf{x}_0}$. Hence,

$$\mathbf{E}_{\mathbf{x}_0 \sim \Pi_{\mathbf{x}_0}} [\mathbf{U}(\mathbf{x}_0)] = \begin{bmatrix} \mathbf{E}_{\mathbf{x}_0 \sim \Pi_{\mathbf{x}_0}} [K_{t_1} U_1(\mathbf{x}_0)] \\ \vdots \\ \mathbf{E}_{\mathbf{x}_0 \sim \Pi_{\mathbf{x}_0}} [K_{t_q} U_q(\mathbf{x}_0)] \end{bmatrix} = \begin{bmatrix} \mathbf{E}_{\mathbf{x}_{t_1} \sim \Phi_*^{t_1} \Pi_{\mathbf{x}_0}} [U_1(\mathbf{x}_{t_1})] \\ \vdots \\ \mathbf{E}_{\mathbf{x}_{t_q} \sim \Phi_*^{t_q} \Pi_{\mathbf{x}_0}} [U_q(\mathbf{x}_{t_q})] \end{bmatrix}.$$

This identity has been used previously in [16,17] to reduce the computational burden of computing the expected values in the right-hand side, which requires sampling from the measures $\Phi_*^{t_1} \Pi_{\mathbf{x}_0}, \dots, \Phi_*^{t_q} \Pi_{\mathbf{x}_0}$, by sampling from $\Pi_{\mathbf{x}_0}$ and then applying the maps $\Phi^{t_1}, \dots, \Phi^{t_q}$.

2.1.3. Parametric systems

In the random differential equation (2) only \mathbf{x}_0 is assumed to be uncertain. However, random differential equations in \mathbb{R}^d of the form

$$\begin{cases} \dot{\mathbf{y}}(t) = \mathbf{g}(t, \mathbf{y}(t), \boldsymbol{\vartheta}_0) \\ \mathbf{y}(0) = \mathbf{y}_0 \end{cases} \quad \text{where } (\mathbf{y}_0, \boldsymbol{\vartheta}_0) \sim \Pi_{(\mathbf{y}_0, \boldsymbol{\vartheta}_0)} \quad (5)$$

where *both* the initial condition \mathbf{y}_0 and a p -dimensional parameter $\boldsymbol{\vartheta}_0$ distribute according to an unknown probability measure $\Pi_{(\mathbf{y}_0, \boldsymbol{\vartheta}_0)}$ are frequent in practice. An equation of this form can be represented as the equation of the form (2) in $\mathbb{R}^d \times \mathbb{R}^p$ by defining

$$\mathbf{x} := \begin{bmatrix} \mathbf{y} \\ \boldsymbol{\vartheta} \end{bmatrix}, \quad \mathbf{x}_0 := \begin{bmatrix} \mathbf{y}_0 \\ \boldsymbol{\vartheta}_0 \end{bmatrix}, \quad \mathbf{f}(t, \mathbf{x}) = \begin{bmatrix} \mathbf{g}(t, \mathbf{y}, \boldsymbol{\vartheta}) \\ 0 \end{bmatrix} \quad \text{and} \quad \Pi_{\mathbf{x}_0} = \Pi_{(\mathbf{y}_0, \boldsymbol{\vartheta}_0)},$$

and then considering the random differential equation

$$\begin{cases} \dot{\mathbf{x}}(t) = \mathbf{f}(t, \mathbf{x}(t)) \\ \mathbf{x}(0) = \mathbf{x}_0 \end{cases} \quad \text{where } \mathbf{x}_0 \sim \Pi_{\mathbf{x}_0}$$

The first d state variables of a trajectory of this system correspond to a trajectory of (5). In this form, the uncertainty is on the initial condition and thus our methods can be applied to quantify the uncertainty in any *finite-dimensional parameter* of the equation.

2.2. Tikhonov regularization as non-parametric maximum a posteriori

To define a family $\{\Pi_{\mathbf{x}_0|\theta} : \theta \in \Theta\}$ for \mathbf{x}_0 that can adapt to complex data we proceed as follows. First, in most practical applications we can usually assume that the initial condition \mathbf{x}_0 must belong to a sufficiently large *bounded* set. Therefore, we choose a compact set $\Omega \subset \mathbb{R}^d$ containing the possible values of \mathbf{x}_0 , and we fix a finite positive Borel measure μ_Ω on Ω that will act as a *dominating measure*; we further assume, without loss of generality, that Ω is the support of μ_Ω . We will restrict ourselves to probability measures on Ω that have a density ρ with respect to μ_Ω , i.e., measures for which we have $\Pi = \rho \mu_\Omega$, and for which the densities belong to the space of (classes of equivalence of μ_Ω -almost everywhere equal) μ_Ω -square-integrable functions

$$L_\mu^2(\Omega) := \left\{ \varphi : \Omega \rightarrow \mathbb{R} : \int_{\Omega} \varphi(\mathbf{x}_0)^2 d\mu_\Omega(\mathbf{x}_0) < \infty \right\},$$

which becomes a Hilbert space when endowed with the inner product

$$\langle \varphi, \psi \rangle_{L_\mu^2} = \int_{\Omega} \varphi(\mathbf{x}_0) \psi(\mathbf{x}_0) d\mu_\Omega(\mathbf{x}_0).$$

To ensure that the integral expressions that arise are well-defined, we assume that the QoIs also belong to this space, whence

$$\int_{\Omega} \|U(\mathbf{x}_0)\|_2^2 d\mu_\Omega(\mathbf{x}_0) < \infty. \quad (6)$$

Since μ_Ω is finite, the functions in $L_\mu^2(\Omega)$ are absolutely integrable and $1 \in L_\mu^2(\Omega)$. As a consequence, the subset of probability densities

$$P_\mu := \left\{ \rho \in L_\mu^2(\Omega) : \rho \geq 0, \int_{\Omega} \rho(\mathbf{x}_0) d\mu_\Omega(\mathbf{x}_0) = 1 \right\}$$

in $L_\mu^2(\Omega)$ is closed and convex.

Restricting ourselves to densities in P_μ allows us to leverage standard techniques to construct finite-dimensional approximations to approximate a density ρ . Let $V_n \subset L_\mu^2(\Omega)$ be a finite dimensional subspace with $\dim(V_n) = n$. Since not all elements of V_n are densities, we define

$$P_{\mu,n} := V_n \cap P_\mu \quad (7)$$

Let $\varphi_1, \dots, \varphi_n$ be an orthonormal basis for V_n and, to simplify notation, define

$$\boldsymbol{\varphi} = \begin{bmatrix} \varphi_1 \\ \vdots \\ \varphi_n \end{bmatrix}.$$

When $\rho \in P_{\mu,n}$ we can represent the density as

$$\rho = \langle \boldsymbol{a}, \boldsymbol{\varphi} \rangle = \sum_{i=1}^n a_i \varphi_i \quad (8)$$

for a suitable coordinate vector $\boldsymbol{a} \in \mathbb{R}^n$. An appealing property of this representation is that the expected value of any function $f \in L_\mu^2(\Omega)$ with respect to this density depends only on the orthogonal projection f_{V_n} of f onto V_n . In fact,

$$\mathbf{E}_{\mathbf{x}_0 \sim \rho}[f(\mathbf{x}_0)] = \sum_{i=1}^n a_i \int_{\Omega} \varphi_i(\mathbf{x}_0) f(\mathbf{x}_0) d\mu_\Omega(\mathbf{x}_0) = \sum_{i=1}^n a_i \int_{\Omega} \varphi_i(\mathbf{x}_0) f_{V_n}(\mathbf{x}_0) d\mu_\Omega(\mathbf{x}_0) = \mathbf{E}_{\mathbf{x}_0 \sim \rho}[f_{V_n}(\mathbf{x}_0)].$$

This becomes useful, for instance, when the distance from f to V_n is small, or when V_n encodes some desirable property of ρ , such as smoothness or localization.

It is not clear *a priori* if $P_{\mu,n}$ is empty. To ensure it is not, and to parameterize densities of the form (8) in terms of their coordinates, we need to first restrict the possible values of this vector to ensure that the expansion defines a density. We defer the proof of this proposition to Appendix A.1.

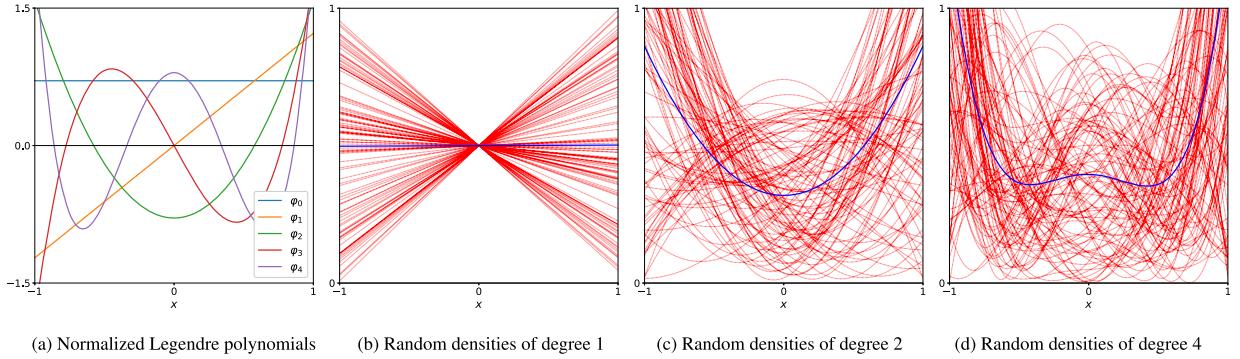


Fig. 1. Random densities on $\Omega = [-1, 1]$ with μ_Ω the Lebesgue measure sampled from the set $P_{\mu,n}$ generated by the normalized Legendre polynomials and the prior in (9). The blue line shows the expected density, whereas the red lines show sampled densities. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

Proposition 1. Suppose that $\varphi_1, \dots, \varphi_n$ are bounded and that φ_1 is strictly positive on Ω . Then there exists a convex set $\Theta_{V_n} \subset \mathbb{R}^{n-1}$ with non-empty interior, a vector $a_n \in \mathbb{R}^n$, and a $n \times (n-1)$ matrix V with $V^*V = I$ and $V^*a_n = 0$ such that

$$\langle a, \varphi \rangle \in P_{\mu,n} \Leftrightarrow \exists \theta \in \Theta_{V_n} : a = a_n + V\theta.$$

In particular, $P_{\mu,n}$ is non-empty.

Proposition 1 and the representation (8) allows us to identify the set of parameters Θ_{V_n} with the set of densities $P_{\mu,n}$ whence a prior on $\theta \in \Theta_{V_n}$ will induce a prior on $\rho \in P_{\mu,n}$. Fix $\lambda > 0$ and define on \mathbb{R}^{n-1} the probability density with respect to the Lebesgue measure given by

$$\rho_\theta(\theta) = \frac{\chi_{\Theta_{V_n}}(\theta) e^{-\lambda \|\theta\|_2^2/2}}{\int_{\Theta_{V_n}} e^{-\lambda \|\theta\|_2^2/2} d\theta}, \quad (9)$$

where the normalizing factor is strictly positive as Θ_{V_n} has non-empty interior, and where $\chi_{\Theta_{V_n}}$ is the indicator function of the set Θ_{V_n} . In Fig. 1 we see how this approach can help us draw random densities from $P_{\mu,n}$.

Suppose that $\Pi_{u|\bar{u}}$ has density

$$\rho_{u|\bar{u}}(u) \propto e^{-\psi(u, \bar{u})}$$

for some function ψ on $\mathbb{R}^d \times \mathbb{R}^d$. If we use the prior for θ in (9) in the model (4) then the likelihood is proportional to

$$\chi_{\Theta_{V_n}}(\theta) \exp\left(-\frac{\|y - u\|_2^2}{2\sigma^2}\right) \exp\left(-\psi\left(u, \int_{\Omega} U(x_0) \rho_\theta(x_0) d\mu_\Omega(x_0)\right)\right) \exp\left(-\frac{\lambda}{2} \|\theta\|_2^2\right)$$

where the constant factors depend on λ and σ . Instead of attempting to compute, or to sample from, the posterior we will compute the maximum *a posteriori* (MAP) estimate for θ , which is the solution to

$$\underset{\theta \in \mathbb{R}^s, u \in \mathbb{R}^q}{\text{maximize}} \quad \log(\chi_{\Theta_{V_n}}(\theta)) - \frac{1}{2\sigma^2} \|y - u\|_2^2 - \psi\left(u, \int_{\Omega} U(x_0) \rho_\theta(x_0) d\mu_\Omega(x_0)\right) - \frac{\lambda}{2} \|\theta\|_2^2$$

The term $\log(\chi_{\Theta_{V_n}}(\theta))$ essentially imposes the constraint $\theta \in \Theta_{V_n}$ whereas u remains unconstrained. Hence, we write

$$\underset{\theta \in \Theta_{V_n}}{\text{maximize}} \quad -L\left(\sum_{i=1}^s \theta_i \int_{\Omega} U(x_0) \varphi_i(x_0) d\mu_\Omega(x_0), y\right) - \frac{\lambda}{2} \|\theta\|_2^2$$

where

$$L(\bar{u}, y) := \inf_{u \in \mathbb{R}^q} \left(\psi(u, \bar{u}) + \frac{1}{2\sigma^2} \|u - y\|_2^2 \right)$$

is the *Moreau envelope* [26, Definition 1.22] of ψ on its first argument. The envelope enjoys some regularity, namely, if $u \mapsto \psi(u, \bar{u})$ is lower semicontinuous and majorizes a quadratic for a fixed \bar{u} then $y \mapsto L(\bar{u}, y)$ is continuous [26, Example 1.24 and Theorem 1.25].

This problem is the classical problem of *ridge regression* or *Tikhonov regularization* to estimate θ . It is clearly a finite-dimensional problem and, when L is convex on \bar{u} , it is also convex. We can leverage Proposition 1 to represent this problem in terms of ρ itself. In fact, since $\varphi_1, \dots, \varphi_n$ are orthonormal, we have the isometry

$$\|\rho\|_{L_\mu^2}^2 = \|\mathbf{a}\|_2^2 = \|\mathbf{a}_n + \mathbf{V}\theta\|_2^2 = \|\mathbf{a}_n\|_2^2 + \|\theta\|_2^2.$$

Since the constraint $\theta \in \Theta_{V_n}$ is equivalent to $\rho_\theta \in P_{\mu,n}$ and the norm of \mathbf{a}_n is independent of θ , finding the MAP for θ is equivalent to finding the *density* that solves the problem

$$\underset{\rho \in P_{\mu,n}}{\text{maximize}} \quad -L \left(\int_{\Omega} \mathbf{U}(\mathbf{x}_0) \rho(\mathbf{x}_0) d\mu_{\Omega}(\mathbf{x}_0), \mathbf{y} \right) - \frac{\lambda}{2} \|\rho\|_{L_\mu^2}^2. \quad (10)$$

This problem is also a form of regularization, but for the density itself, which emphasizes the choice of V_n over the particular choice of basis $\varphi_1, \dots, \varphi_n$. If $\{V_n\}_{n \in \mathbb{N}}$ is an increasing sequence of subspaces for which

$$\overline{\bigcup_{n \in \mathbb{N}} V_n} = L_\mu^2(\Omega) \quad (11)$$

then, from (7), we conclude that the asymptotic regime when $n \gg 1$ leads to the *infinite-dimensional problem*

$$\underset{\rho \in P_\mu}{\text{maximize}} \quad -L \left(\int_{\Omega} \mathbf{U}(\mathbf{x}_0) \rho(\mathbf{x}_0) d\mu_{\Omega}(\mathbf{x}_0), \mathbf{y} \right) - \frac{\lambda}{2} \|\rho\|_{L_\mu^2}^2. \quad (12)$$

This is a *nonparametric method* in that the parameter, i.e., the density ρ , lies on an infinite-dimensional space. We can prove that, under suitable conditions and in a suitable sense, the sequence of MAP estimates found by solving (10) for increasingly large values of n converges to a density found by solving (12). We defer the proof of the proposition to A.2

Proposition 2. Suppose that $\{V_n\}_{n \in \mathbb{N}}$ is an increasing sequence of finite-dimensional subspaces of $L_\mu^2(\Omega)$ for which (11) holds, and suppose L is non-negative and continuous on its first argument. Furthermore, suppose that for every $n \in \mathbb{N}$ the problem (10) has an optimal solution $\rho^{\star,n}$ and that (12) has at least one solution. Then a subsequence of $\{\rho^{\star,n}\}_{n \in \mathbb{N}}$ converges weakly to a solution to (12).

The proposition formalizes the intuition that a solution to (12) represents the limit of MAP estimates for the density. However, to be truly considered a MAP estimate, we would need to introduce a suitable prior on the set P_μ . This would result in a Bayesian nonparametric method [23]. We do not address the question whether there exists such a prior. Instead, we focus on analyzing the properties of (12), that is, the existence, uniqueness, and features of its optimal solution.

3. Analysis of the variational problem

We study the family of problems induced by (12) given by

$$\begin{aligned} & \underset{\rho \in L_\mu^2(\Omega), \bar{u} \in C: \rho \geq 0}{\text{minimize}} \quad \frac{1}{2} \|\rho\|_{L_\mu^2}^2 + \frac{1}{\lambda} L(\mathbf{A}\bar{u}, \mathbf{y}) \\ & \text{subject to} \quad \int_{\Omega} \rho(\mathbf{x}_0) d\mu_{\Omega}(\mathbf{x}_0) = 1, \quad \int_{\Omega} \rho(\mathbf{x}_0) \mathbf{U}(\mathbf{x}_0) d\mu_{\Omega}(\mathbf{x}_0) = \bar{u}, \end{aligned} \quad (13)$$

where $C \subset \mathbb{R}^q$ is a closed convex set representing restrictions on \bar{u} induced by the prior $\Pi_{\bar{u}|\bar{u}}$, e.g., non-negativity constraints, bound constraints or norm constraints, and L is a loss function which we assume is convex on its first argument. The regularization parameter $\lambda > 0$ controls the trade-off between the loss and the squared-norm penalization term on ρ .

Proposition 3. There exists a unique minimizer to (13).

We defer the proof of this proposition to Appendix B.1. Observe that the method proposed in [16] is equivalent to solving

$$\begin{aligned} & \underset{\rho \in L_\mu^2(\Omega): \rho \geq 0}{\text{minimize}} \quad \frac{1}{2} \|\rho\|_{L_\mu^2}^2 \\ & \text{subject to} \quad \int_{\Omega} \rho(\mathbf{x}_0) d\mu_{\Omega}(\mathbf{x}_0) = 1, \quad \int_{\Omega} \rho(\mathbf{x}_0) \mathbf{U}(\mathbf{x}_0) d\mu_{\Omega}(\mathbf{x}_0) = \bar{u} \end{aligned} \quad (14)$$

and that this is a particular case of (13).

From now on, we let ρ^* be the optimal solution to (13) and we let $\Pi^* := \rho^* \mu_{\Omega}$ be the probability measure associated to the optimal density ρ^* . Since (13) is an infinite-dimensional problem, it is necessary to find numerical approximations to its solution. Although this may be achieved by constructing suitable spaces V_n , we leverage convex duality instead to reformulate the problem in an equivalent form that is amenable to accurate approximations of its solution.

3.1. The dual problem

Instead of attempting to solve or to approximate (13) directly, we leverage convex duality to find an equivalent formulation that is amenable to numerical methods for high-dimensional optimization problems. We briefly review the arguments behind convex duality. First, rewrite (13) as the equivalent problem,

$$\begin{aligned} & \underset{\rho \in L_\mu^2(\Omega), \bar{u} \in \mathbb{R}^q, v \in \mathbb{R}^m}{\text{minimize}} \quad \frac{1}{2} \|\rho\|_{L_\mu^2}^2 + \mathbb{I}_{L_\mu^{2+}}(\rho) + \frac{1}{\lambda} L(v, y) + \mathbb{I}_C(\bar{u}) \\ & \text{subject to} \quad \int_{\Omega} \rho(x_0) d\mu_{\Omega}(x_0) = 1, \quad \int_{\Omega} \rho(x_0) U(x_0) d\mu_{\Omega}(x_0) = \bar{u}, \quad A\bar{u} = v, \end{aligned} \quad (15)$$

where $\mathbb{I}_{L_\mu^{2+}}$ is the indicator function of the closed and convex cone of non-negative functions in $L_\mu^2(\Omega)$, and \mathbb{I}_C is the indicator function of the convex set $C \subset \mathbb{R}^q$. If we define the Hilbert spaces

$$X_P := L_\mu^2(\Omega) \times \mathbb{R}^q \times \mathbb{R}^m \quad \text{and} \quad X_D := \mathbb{R} \times \mathbb{R}^q \times \mathbb{R}^m,$$

the linear map $\mathcal{B} : X_P \rightarrow X_D$ as

$$\mathcal{B}(\rho, \bar{u}, v) = \begin{bmatrix} -\int_{\Omega} \rho(x_0) d\mu_{\Omega}(x_0) \\ -\int_{\Omega} \rho(x_0) U(x_0) d\mu_{\Omega}(x_0) + \bar{u} \\ -A\bar{u} + v \end{bmatrix},$$

the vector $r \in X_D$ as

$$r = \begin{bmatrix} -1 \\ 0 \\ 0 \end{bmatrix},$$

and the convex functions $f : X_P \rightarrow \mathbb{R}$ and $g : X_D \rightarrow \mathbb{R}$ as

$$f(\rho, \bar{u}, v) = \frac{1}{2} \|\rho\|_{L_\mu^2}^2 + \mathbb{I}_{L_\mu^{2+}}(\rho) + \frac{1}{\lambda} L(v, y) + \mathbb{I}_C(\bar{u}),$$

$$g(\eta, \omega, \xi) = \mathbb{I}_0((\eta, \omega, \xi) - r),$$

then (15) can be written equivalently as the unconstrained problem

$$\underset{(\rho, \bar{u}, v) \in X_P}{\text{minimize}} \quad f(\rho, \bar{u}, v) + g(\mathcal{B}(\rho, \bar{u}, v) - r). \quad (16)$$

We call this the *primal problem* and its optimal solution (ρ^*, \bar{u}^*, v^*) the *optimal primal variables*. Its *dual problem* is

$$\underset{(\eta, \omega, \xi) \in X_D}{\text{maximize}} \quad -f^*(-\mathcal{B}^*(\eta, \omega, \xi)) - g^*(\eta, \omega, \xi) \quad (17)$$

where f^* and g^* are the convex conjugates [27, Definition 13.1] of f and g and $\mathcal{B}^* : X_D \rightarrow X_P$ is the adjoint of \mathcal{B} . The optimal solution $(\eta^*, \omega^*, \xi^*)$ yields the *optimal dual variables*. In some cases, solving the dual problem may be much easier than solving the primal problem. However, this is useful only if we are able to recover the optimal primal variables by solving the dual problem. When strong duality holds, the optimal values for (16) and (17) coincide, and we may obtain the optimal primal variables from the optimal dual variables. In the literature, this is also referred as *total duality*. We defer the proof of the following theorem to Appendix B.2.

Theorem 1. *Strong duality holds for (13). The dual problem is the finite-dimensional problem*

$$\underset{\eta, \omega, \xi}{\text{maximize}} \quad -H(\eta, \omega) + \eta - \frac{1}{\lambda} L^*(-\lambda \xi, y) - \mathbb{I}_C^*(A^t \xi - \omega),$$

where L^* is the convex dual of L on its first variable, \mathbb{I}_C^* is the convex dual of \mathbb{I}_C and $H : \mathbb{R} \times \mathbb{R}^q \rightarrow \mathbb{R}$ is the convex function

$$H(\alpha, \beta) = \frac{1}{2} \int_{\Omega} (\alpha + \langle \beta, U(x_0) \rangle_2)_+^2 d\mu_{\Omega}(x_0). \quad (18)$$

Furthermore, if $(\eta^*, \omega^*, \xi^*)$ are the optimal dual variables then the optimal primal variables are

$$\begin{aligned} \rho^*(x_0) &:= (\eta^* + \langle \omega^*, U(x_0) \rangle_2)_+, \\ \bar{u}^* &:= \int_{\Omega} (\eta^* + \langle \omega^*, U(x_0) \rangle_2)_+ U(x_0) d\mu_{\Omega}(x_0), \\ v^* &:= A\bar{u}^*. \end{aligned}$$

There are several consequences of this result. First, the theorem fully characterizes class of density functions that can be obtained by solving (13). Hence, when the optimal dual variables η^*, ω^* are known, we may evaluate the optimal probability density to compute any statistics about the initial condition and about the QoIs. Second, the class of density functions is no more complex than that defined by the QoIs. To be concrete, if U_1, \dots, U_q are the components of U , then the optimal probability density can only have the form

$$\rho^*(\mathbf{x}_0) = \left(\eta^* + \sum_{i=1}^p \omega_i^* U_i(\mathbf{x}_0) \right)_+$$

or, equivalently, it must be the positive part of some function in the span of $1, U_1, \dots, U_p \in L_\mu^2(\Omega)$. Third, we may change the properties of the optimal density found by solving (13) by modifying the reference measure μ_Ω . Fourth, it shows that we can solve (13) by solving its dual problem, which is equivalent to the convex minimization problem

$$\underset{\eta, \omega, \xi}{\text{minimize}} \quad H(\eta, \omega) - \eta + \frac{1}{\lambda} L^*(-\lambda \xi, \mathbf{y}) + \mathbb{I}_C^*(\mathbf{A}' \xi - \omega). \quad (19)$$

This problem is not only finite-dimensional, having $1 + q + m$ variables, but also unconstrained. Each one of the terms in the dual objective encodes the information about the problem: the first about the QoI U , the second about the loss L , and the third about the *a priori* constraints we impose on \bar{u} . Usually, the last two terms reflect modeling choices and there is some freedom to select them. In particular, we may choose the priors in such a way that the loss L and the set of constraints C have a small computational burden when solving (19). In contrast, the first term depends on the QoIs and the available data. As a consequence, it is typically the first term that will contribute most to the computational cost.

The complexity of solving an unconstrained optimization problem depends on the regularity of its objective function. An important property of H is that it is differentiable, with Lipschitz continuous derivatives. We defer the proof to Appendix B.4.

Proposition 4. *The function H defined in (18) is continuously differentiable, with partial derivatives*

$$\begin{aligned} \partial_\alpha H(\alpha, \beta) &= \int_{\Omega} (\alpha + \langle \beta, U(\mathbf{x}_0) \rangle_2)_+ d\mu_\Omega(\mathbf{x}_0), \\ \nabla_\beta H(\alpha, \beta) &= \int_{\Omega} U(\mathbf{x}_0) (\alpha + \langle \beta, U(\mathbf{x}_0) \rangle_2)_+ d\mu_\Omega(\mathbf{x}_0). \end{aligned}$$

In particular, ∇H is Lipschitz continuous with

$$\text{Lip}(\nabla H) \leq \int_{\Omega} \max\{1, \|U(\mathbf{x}_0)\|_2^2\} d\mu_\Omega(\mathbf{x}_0).$$

Therefore, when L^* and \mathbb{I}_C^* are differentiable, we may solve (19) using gradient descent or projected gradient descent. These algorithms are simple to implement, and their accelerated versions attain optimal convergence rates on the class of convex functions.

A natural question is whether we can leverage similar arguments to solve (13) when ρ is restricted to a finite-dimensional subspace V_n of $L_\mu^2(\Omega)$. Interestingly, a similar result to Theorem 1 holds in this case.

Theorem 2. *Let $V_n \subset L_\mu^2(\Omega)$ be a finite-dimensional subspace containing at least one strictly positive function. Let $1_{V_n}, U_{1,V_n}, \dots, U_{q,V_n}$ be the orthogonal projections of $1, U_1, \dots, U_q$ onto V_n and define U_{V_n} accordingly. Consider the problem*

$$\begin{aligned} \underset{\rho \in V_n, \bar{u} \in \mathbb{R}^q, \mathbf{v} \in \mathbb{R}^m}{\text{minimize}} \quad & \frac{1}{2} \|\rho\|_{L_\mu^2}^2 + \mathbb{I}_{V_n^+}(\rho) + \frac{1}{\lambda} L(\mathbf{v}, \mathbf{y}) + \mathbb{I}_C(\bar{u}) \\ \text{subject to} \quad & \int_{\Omega} \rho(\mathbf{x}_0) d\mu_\Omega(\mathbf{x}_0) = 1, \quad \int_{\Omega} \rho(\mathbf{x}_0) U(\mathbf{x}_0) d\mu_\Omega(\mathbf{x}_0) = \bar{u}, \quad \mathbf{A}\bar{u} = \mathbf{v}. \end{aligned}$$

where V_n^+ is the closed convex cone of non-negative functions in V_n . Then strong duality holds for this problem, and its dual problem is the finite-dimensional problem

$$\underset{\eta, \omega, \xi}{\text{maximize}} \quad -H_n(\eta, \omega) + \eta - \frac{1}{\lambda} L^*(-\lambda \xi, \mathbf{y}) - \mathbb{I}_C^*(\mathbf{A}' \xi - \omega).$$

where $H_n : \mathbb{R} \times \mathbb{R}^q \rightarrow \mathbb{R}$ is the convex function

$$H_n(\alpha, \beta) = \frac{1}{2} \int_{\Omega} \left(\alpha 1_{V_n}(\mathbf{x}_0) + \langle \beta, U_{V_n}(\mathbf{x}_0) \rangle_2 \right)_+^2 d\mu_\Omega(\mathbf{x}_0).$$

Furthermore, if the primal problem is feasible, and $(\eta^*, \omega^*, \xi^*)$ are the optimal dual variables then the optimal primal variables are

$$\begin{aligned}\rho^*(\mathbf{x}_0) &:= \left(\eta^* 1_{V_n}(\mathbf{x}_0) + \langle \boldsymbol{\omega}^*, \mathbf{U}_{V_n}(\mathbf{x}_0) \rangle_2 \right)_+, \\ \bar{\mathbf{u}}^* &:= \int_{\Omega} \left(\eta^* 1_{V_n}(\mathbf{x}_0) + \langle \boldsymbol{\omega}^*, \mathbf{U}_{V_n}(\mathbf{x}_0) \rangle_2 \right)_+ \mathbf{U}(\mathbf{x}_0) d\mu_{\Omega}(\mathbf{x}_0), \\ \mathbf{v}^* &:= \mathbf{A}\bar{\mathbf{u}}^*.\end{aligned}$$

3.2. Covariance terms

A critical condition to apply duality for (13) is that the constraints on the expectation are linear in the density, i.e.,

$$\int_{\Omega} \mathbf{U}(\mathbf{x}_0) \rho(\mathbf{x}_0) d\mu_{\Omega}(\mathbf{x}_0) = \bar{\mathbf{u}}.$$

In general, this constraint could be replaced by any other convex constraint and our arguments would apply with minor modifications. An exception would be any explicit constraints on the covariance matrix, as these would take the form

$$\boldsymbol{\mu} := \int_{\Omega} \mathbf{x}_0 \rho(\mathbf{x}_0) d\mu_{\Omega}(\mathbf{x}_0) \quad \text{and} \quad \boldsymbol{\Sigma} := \int_{\Omega} (\mathbf{x}_0 - \boldsymbol{\mu})(\mathbf{x}_0 - \boldsymbol{\mu})^T \rho(\mathbf{x}_0) d\mu_{\Omega}(\mathbf{x}_0);$$

while the former is convex, the second is not. We can overcome this limitation as follows. From the linear equality constraints

$$\boldsymbol{\mu} := \int_{\Omega} \mathbf{x}_0 \rho(\mathbf{x}_0) d\mu_{\Omega}(\mathbf{x}_0) \quad \text{and} \quad \mathbf{X} := \int_{\Omega} \mathbf{x}_0 \mathbf{x}_0^T \rho(\mathbf{x}_0) d\mu_{\Omega}(\mathbf{x}_0)$$

we recover the covariance as $\boldsymbol{\Sigma} = \mathbf{X} - \boldsymbol{\mu}\boldsymbol{\mu}^T$. Then, a convex loss or a convex penalization term for \mathbf{X} and $\boldsymbol{\mu}$ can be added to the objective function in (13). Similarly, some constraints can be imposed by choosing a suitable set C . In both cases, the same arguments used in Theorem 1 allow us to compute the dual of the problem.

3.3. Special problem instances

3.3.1. Quadratic loss

A typical choice for the loss in problems of the form (13) is the *quadratic loss*. In this case,

$$L(\mathbf{v}, \mathbf{y}) = \frac{1}{2} \|\mathbf{v} - \mathbf{y}\|_2^2 \quad \text{and} \quad C = \mathbb{R}^q.$$

This loss is obtained, for instance, when using the prior

$$\rho_{\mathbf{u} \mid \bar{\mathbf{u}}}(\mathbf{u}) \propto e^{-\frac{1}{2\tau^2} \|\mathbf{u} - \bar{\mathbf{u}}\|_2^2}.$$

From this,

$$L^*(\mathbf{v}, \mathbf{y}) = \frac{1}{2} \|\xi\|_2^2 + \langle \xi, \mathbf{y} \rangle \quad \text{and} \quad \mathbb{I}_C^*(\boldsymbol{\omega}) = \mathbb{I}_0(\boldsymbol{\omega}).$$

Since

$$\mathbb{I}_0(\mathbf{A}'\xi - \boldsymbol{\omega}) < \infty \Leftrightarrow \boldsymbol{\omega} = \mathbf{A}'\xi,$$

this instance of the problem becomes

$$\underset{\eta, \xi}{\text{minimize}} \quad H(\eta, \mathbf{A}'\xi) - \eta + \frac{\lambda}{2} \|\xi\|_2^2 - \langle \xi, \mathbf{y} \rangle_2.$$

This instance of (23) is unconstrained, and the objective is differentiable with Lipschitz gradient. Hence, it can be solved using gradient descent with acceleration [28–30]. The computational burden at each iteration of the algorithm is then dominated by the computing the value of H and of its derivatives.

3.3.2. Lebesgue measure and star domains

In practice, the computational burden of solving (19) will be dominated by the cost of evaluating H and its derivatives. However, for some special choices of Ω and μ_{Ω} we can represent H as an integral over the closed Euclidean ball \bar{B}_2 . In this case, instead of developing quadrature rules for every Ω we may use quadrature methods over the fixed set \bar{B}_2 ; the trade-off is that the regularity of the integrand may decrease.

Recall that a closed set $\Omega \subset \mathbb{R}^d$ is *star-shaped* with respect to some $\mathbf{x}_c \in \Omega$ if it contains a neighborhood of \mathbf{x}_c and if for any $\mathbf{x} \in \Omega$ the ray

$$\{\mathbf{x}_c + s(\mathbf{x} - \mathbf{x}_c) : s > 0\}$$

does not intersect the boundary of Ω more than once [31, Section 2]. Its Minkowski functional

$$\gamma_0(\mathbf{x}) := \inf \{t > 0 : t\mathbf{x} \in \Omega_0\},$$

is positively homogeneous, non-negative and continuous and [31, Theorem 1]

$$\Omega := \{\mathbf{x} \in \mathbb{R}^q : \gamma_0(\mathbf{x}) \leq 1\}.$$

Suppose that Ω is star-shaped with respect to some

$$\mathbf{x}_c \in \text{int}(\Omega).$$

By translating Ω we have

$$H(\alpha, \beta) = \frac{1}{2} \int_{\Omega_0} (\alpha + \langle \beta, \mathbf{U}(\mathbf{x}_0 + \mathbf{x}_c) \rangle_2)_+^2 d\mathbf{x}_0$$

where $\Omega_0 := \Omega - \mathbf{x}_c$ contains a neighborhood of the origin. Let γ_0 be the Minkowski functional associated to Ω_0 and let

$$T(\mathbf{x}) = \begin{cases} \frac{\|\mathbf{x}\|_2}{\gamma_0(\mathbf{x})} \mathbf{x} & \mathbf{x} \neq 0 \\ \mathbf{0} & \mathbf{x} = 0. \end{cases} \quad (20)$$

Then

$$T : \bar{B}_2 \rightarrow \Omega_0$$

is an homeomorphism. First, T is continuous by inspection and

$$\gamma_0(T(\mathbf{x})) = \|\mathbf{x}\|_2$$

implies $T : \bar{B}_2 \rightarrow \Omega_0$. Its inverse is

$$T^{-1}(\mathbf{x}) = \frac{\gamma_0(\mathbf{x})}{\|\mathbf{x}\|_2} \mathbf{x}. \quad (21)$$

Since the same arguments show $T^{-1} : \Omega_0 \rightarrow \bar{B}_2$ is continuous we conclude T is an homeomorphism. This map allows us to represent H as an integral over the closed unit ball by changing the integrand. We defer the proof of the following result to Appendix B.5.

Theorem 3. *Let Ω be a star shaped set with respect to some $\mathbf{x}_c \in \text{int}(\Omega)$ and let μ_Ω be the Lebesgue measure. Under the change of variables $\mathbf{x} \mapsto T(\mathbf{x})$ we have*

$$H(\alpha, \beta) = \frac{1}{2} \int_{\bar{B}_2} \left(\alpha + \left\langle \beta, \mathbf{U} \left(\frac{\|\mathbf{x}_0\|_2}{\gamma_0(\mathbf{x}_0)} \mathbf{x}_0 + \mathbf{x}_c \right) \right\rangle_2 \right)_+^2 \left(\frac{\|\mathbf{x}_0\|_2}{\gamma_0(\mathbf{x}_0)} \right)^d d\mathbf{x}_0. \quad (22)$$

The representation in (22) nicely shows the influence of \mathbf{U} and Ω on the dual objective. The integrand is a product of two factors. The first one is essentially the same as the integrand in (18). The second is a weight that quantifies how different is Ω_0 from \bar{B}_2 . This weight is homogeneous of degree 0 and it depends only on the *direction* of the integrand, but not on its magnitude. Remark the dependence on the quotient is *exponential* on the dimension, whence the points that will contribute the most to the integral are those where the quotient is larger.

4. Interpretation as regularized regression

The problem (13) can be interpreted both as using MAP to estimate a density ρ , or as a regression problem to estimate the expected value \bar{u} . In fact, if we let $R_U = R_U(\bar{u})$ be the optimal value to (14) then (13) becomes

$$\underset{\bar{u} \in C}{\text{minimize}} \quad R_U(\bar{u}) + \frac{1}{\lambda} L(\mathbf{A}\bar{u}, \mathbf{y}). \quad (23)$$

This is a *regularized regression problem* to estimate \bar{u} . Observe that it is finite-dimensional, and that the number of variables is equal to the number q of QoIs. The terms that depend on the density ρ in (13) act implicitly as a regularizer on the estimate of \bar{u} . In this context, it is of interest to determine the properties of the regularizer. We defer the proof of the proposition to Appendix C.1.

Proposition 5. *Suppose*

$$\int_{\Omega} \|\mathbf{U}(\mathbf{x}_0)\|_2^2 d\mu_\Omega(\mathbf{x}_0) < \infty.$$

Then, the function $R_U : \mathbb{R}^q \rightarrow \mathbb{R}$ is proper and convex. Furthermore, its domain is the convex set

$$\mathbf{dom}(R_U) := \left\{ \int_{\Omega} \rho(\mathbf{x}_0) U(\mathbf{x}_0) d\mu_{\Omega}(\mathbf{x}_0) : \rho \in L^2_{\mu}(\Omega), \rho \geq 0, \int_{\Omega} \rho(\mathbf{x}_0) d\mu_{\Omega}(\mathbf{x}_0) = 1 \right\} \quad (24)$$

Under mild assumptions (23) has at least one optimal solution [27, Proposition 11.13]. One important consequence of this result is that (13) implicitly enforces the constraint $\bar{\mathbf{u}} \in \mathbf{dom}(R_U)$. In turn, this constraint depends both on the choice of \mathbf{U} and the set Ω . In fact,

$$\mathbf{dom}(R_U) \subset \mathbf{cvx}\{U(\mathbf{x}_0) : \mathbf{x}_0 \in \Omega\}.$$

In the special case when Ω is a convex body, we can show that at least in some special cases $\mathbf{dom}(R_U)$ is open and hence this inclusion may be strict. We defer the proof to Appendix C.2.

Proposition 6. Suppose that Ω is a convex body, that μ_{Ω} is the Lebesgue measure, and that

$$U(\mathbf{x}_0) = \mathbf{x}_0.$$

Then $\mathbf{dom}(R_U) = \mathbf{int}(\Omega)$. Furthermore, $R_U(\bar{\mathbf{u}}) \rightarrow \infty$ as $\bar{\mathbf{u}} \rightarrow \mathbf{bd}(\Omega)$.

The statement extends with minor modifications to other choices of measure μ_{Ω} on Ω . Remark that, in general, if μ_{Ω} is absolutely continuous with respect to the Lebesgue measure then the same arguments in the proposition yield

$$\mathbf{dom}(R_U) \subset \mathbf{supp}\left(\frac{d\mu_{\Omega}}{d\mathbf{x}_0}\right).$$

The formulation (23) shows a clear distinction between the full and partial measurement models. The partial measurement model allows us to codify the effect of \mathbf{U} in the regularization term, whereas the measurement matrix \mathbf{A} appears only in the loss. There is an extensive literature on the performance of estimators based on regularized regression. Results of this type can be applied to understand how close is the optimal solution to (23) to the true expectation $\bar{\mathbf{u}}_0$.

4.1. Linear systems and convex sets

When the system (2) is linear and the QoIs are also linear in the state, we obtain a special case of (23). By a slight abuse of notation, we write

$$U(\mathbf{x}_0) = \mathbf{U}\mathbf{x}_0$$

for a $q \times d$ matrix \mathbf{U} . Suppose that Ω is a convex body. Then, we may decompose the regularization function R_U as follows. Let $R_U^{\circ} = R_U^{\circ}(\boldsymbol{\mu})$ be the optimal value of the problem

$$\begin{aligned} & \underset{\rho \in L^2_{\mu}(\Omega) : \rho \geq 0}{\text{minimize}} \quad \frac{1}{2} \int_{\Omega} \rho(\mathbf{x}_0)^2 d\mu_{\Omega}(\mathbf{x}_0) \\ & \text{subject to} \quad \int_{\Omega} \rho(\mathbf{x}_0) d\mu_{\Omega}(\mathbf{x}_0) = 1, \quad \int_{\Omega} \rho(\mathbf{x}_0) \mathbf{x}_0 d\mu_{\Omega}(\mathbf{x}_0) = \boldsymbol{\mu}. \end{aligned}$$

By Proposition 5 and 6 we deduce that R_U° is proper, convex and $\mathbf{dom}(R_U^{\circ}) = \mathbf{int}(\Omega)$. Then, the value of $R_U = R_U(\bar{\mathbf{u}})$ is the optimal value of the problem

$$\underset{\boldsymbol{\mu} \in \mathbb{R}^d}{\text{minimize}} \quad R_U^{\circ}(\boldsymbol{\mu}) \quad \text{subject to} \quad \mathbf{U}\boldsymbol{\mu} = \bar{\mathbf{u}}.$$

Therefore, we can rewrite (23) equivalently as

$$\underset{\boldsymbol{\mu} \in \mathbb{R}^d, \bar{\mathbf{u}} \in C}{\text{minimize}} \quad R_U^{\circ}(\boldsymbol{\mu}) + \frac{1}{\lambda} L(\mathbf{A}\bar{\mathbf{u}}, \mathbf{y}) \quad \text{subject to} \quad \mathbf{U}\boldsymbol{\mu} = \bar{\mathbf{u}}. \quad (25)$$

In some cases, this problem can be further reduced to a standard form. We defer the proof to Appendix C.3.

Proposition 7. Suppose that Ω is a closed Euclidean ball of radius $r > 0$ and center $\boldsymbol{\mu}_0$ and that μ_{Ω} is the Lebesgue measure. Then there exists a non-decreasing convex function $\varphi : [0, 1] \rightarrow \mathbb{R}$ such that

$$R_U^{\circ}(\boldsymbol{\mu}) = \varphi\left(\frac{\|\boldsymbol{\mu} - \boldsymbol{\mu}_0\|_2}{r}\right)$$

and $\varphi(s) \rightarrow \infty$ as $s \rightarrow 1$. Furthermore, (25) becomes

$$\underset{\mu \in \mathbb{R}^d, \bar{u} \in C, s \geq 0}{\text{minimize}} \quad \varphi(s) + \frac{1}{\lambda} L(\mathbf{A}\bar{u}, \mathbf{y}) \quad \text{subject to} \quad \mathbf{U}\mu = \bar{u}, \quad \frac{\|\mu - \mu_0\|_2}{r} \leq s.$$

Therefore, for linear systems and observables, and in the Euclidean ball and (a scaling of) the Lebesgue measure, the regularization effect on the expectation is very similar to that of Tikhonov regularization. This once again shows the implicit effect that the choice of Ω and of reference measure μ_Ω has on the behavior of the method.

5. Implementation

We now discuss some methods that can be used to solve numerically a special instance of (19). We will assume that L^* is differentiable with Lipschitz gradient, that $C = \mathbb{R}^q$ and that the measure μ_Ω is a probability measure on Ω ; to make this assumption explicit, from now on we write Π_Ω . From these assumptions, it follows from the arguments in Section 3.3.1 that the dual problem has the form

$$\underset{\eta, \xi}{\text{minimize}} \quad H(\eta, \mathbf{A}^t \xi) - \eta + \frac{1}{\lambda} L^*(\lambda \xi, \mathbf{y}) \quad (26)$$

Furthermore, if we define

$$h(\alpha, \beta, \mathbf{x}_0) := \frac{1}{2}(\alpha + \langle \beta, \mathbf{U}(\mathbf{x}_0) \rangle)_+$$

then H can be represented as

$$H(\alpha, \beta) = \int_{\Omega} h(\alpha, \beta, \mathbf{x}_0) d\Pi_\Omega(\mathbf{x}_0) = \mathbb{E}_{\mathbf{x}_0 \sim \Pi_\Omega} [h(\alpha, \beta, \mathbf{x}_0)].$$

Therefore, we can interpret (26) as either a deterministic or stochastic convex optimization problem. Depending on this choice, different algorithms can be implemented in practice.

5.1. Gradient descent methods

As (26) is unconstrained with a convex and smooth objective function, perhaps the simplest and most efficient method to solve this problem is accelerated gradient descent (GD) [28,32,29,30]. However, to simplify the exposition we will focus on gradient descent without acceleration. Given (η_0, ξ_0) and a sequence of steps $\{s_n\}_{n \in \mathbb{N}}$ we define for $n \geq 0$ the iterates

$$\begin{aligned} \eta_{n+1} &= \eta_n - s_n(\partial_\alpha H(\eta_n, \mathbf{A}^t \xi_n) - 1), \\ \xi_{n+1} &= \xi_n - s_n(\mathbf{A} \nabla_\beta H(\eta_n, \mathbf{A}^t \xi_n) - \mathbf{y} + \lambda \xi_n), \end{aligned}$$

In practice, the steps can be selected using backtracking, or they can be constant. In this case, from Proposition 4, the step s must satisfy

$$\frac{1}{s} > \frac{1}{2}(1 + \|\mathbf{A}\|_{\text{op}}) \int_{\Omega} \max\{1, \|\mathbf{U}(\mathbf{x}_0)\|_2^2\} d\Pi_\Omega(\mathbf{x}_0) + \frac{1}{2}\lambda \text{Lip}(\nabla L^*)$$

to ensure convergence.

Gradient descent requires evaluating at least the gradient of H at each iteration. Both H and ∇H are defined in terms of integrals that need to be approximated numerically for possibly large d . We discuss cubature formulae and the *sample average approximation* (SAA). Both methods essentially construct an approximation \hat{H} of H from which \hat{H} and $\nabla \hat{H}$ can be computed efficiently. In practice, one can construct a low-accuracy approximation \hat{H} using few nodes or samples, to first find an approximation $(\hat{\eta}^*, \hat{\xi}^*)$ to the solution to (26) in a few iterations, and then use this approximation as an initial iterate to solve (26) using a more accurate approximation.

5.1.1. Cubature formulae

A first approach is to approximate H in low dimensions is to use cubature formulae. They can be designed to exploit the structure of the integrand to obtain theoretical guarantees about their approximation properties [33,34]. A disadvantage is that they need to be tailored to the class of functions over which one wants to achieve an exact approximation, i.e., the *degree* of the formula, and to the domain of integration, e.g., the hypercube [35,36], the simplex [36,37] or the ball [36]. Their complexity, i.e., the number of nodes required for a desired accuracy, increases exponentially as the dimension increases (however, see [35,38]). Furthermore, their accuracy may vary according to the choice of Π_Ω .

When Π_Ω is a scalar multiple of the Lebesgue measure on Ω we can leverage Theorem 3 to represent H as the integral over the unit ball. In this case, we can use cubature formulae for the Euclidean ball. Let

$$h^\circ(\alpha, \beta, \mathbf{x}_0) = \frac{1}{2} \left(\alpha + \left\langle \beta, \mathbf{U} \left(\frac{\|\mathbf{x}_0\|_2}{\gamma_0(\mathbf{x}_0)} \mathbf{x}_0 + \mathbf{x}_c \right) \right\rangle_2 \right)_+^2 \left(\frac{\|\mathbf{x}_0\|_2}{\gamma_0(\mathbf{x}_0)} \right)^d.$$

Theorem 3 allows us to write

$$\begin{aligned} H(\alpha, \beta) &= \frac{|\partial \bar{B}_2(0, 1)|}{|\Omega|} \int_0^1 \left(\frac{1}{|\partial B_2(0, 1)|} \int_{\partial \bar{B}_2(0, 1)} h^\circ(\alpha, \beta, r\omega) dS(\omega) \right) r^{d-1} dr \\ &=: \frac{|\partial \bar{B}_2(0, 1)|}{|\Omega|} \int_0^1 \bar{h}^\circ(\alpha, \beta, r) r^{d-1} dr \end{aligned}$$

Since there are terms in the definition of h° that are homogeneous of degree 0, we have

$$h^\circ(\alpha, \beta, r\omega) = \frac{1}{2} \left(\alpha + \left\langle \beta, U \left(\frac{\|\omega\|_2}{\gamma_0(\omega)} r\omega + x_c \right) \right\rangle_2 \right)_+^2 \left(\frac{\|\omega\|_2}{\gamma_0(\omega)} \right)^d,$$

whence the regularity of $\bar{h}^\circ(\alpha, \beta, r)$ is controlled by that of U . Standard cubature formulae can be used to approximate the integral on the radial variable, e.g., Gauss-Jacobi formulae for $\alpha = d - 1$ and $\beta = 0$. To evaluate \bar{h}° at the nodes, we may use cubature formulae on the unit sphere, e.g., see [39,40].

5.1.2. Sample average approximation

Although cubature formulae provide reasonable approximations, their limitations make them effective only in very low dimensions. As the systems under study become high-dimensional systems, Monte Carlo (MC) and Quasi-Monte Carlo (QMC) methods become more appropriate. In MC the nodes are sampled according to the probability measure Π_Ω whereas in QMC they are chosen from a *low-discrepancy* sequence [41–43]; in both, the nodes are assigned the same weight. In the optimization literature, e.g., see [44], this approach leads to the sample average approximation (SAA) method to solve (26).

In MC the key idea is to use the approximation

$$\hat{H}(\alpha, \beta) := \frac{1}{N} \sum_{i=1}^N h(\alpha, \beta, x_0^{(i)}) \quad (27)$$

where $x_0^{(1)}, \dots, x_0^{(N)}$ are i.i.d. samples from Π_Ω . Remark the samples are not modified once drawn. The gradient can be readily computed as

$$\begin{aligned} \partial_\alpha \hat{H}(\alpha, \beta) &= \frac{1}{N} \sum_{i=1}^N (\alpha + \langle \beta, U(x_0^{(i)}) \rangle_2)_+, \\ \nabla_\beta \hat{H}(\alpha, \beta) &= \frac{1}{N} \sum_{i=1}^N (\alpha + \langle \beta, U(x_0^{(i)}) \rangle_2)_+ U(x_0^{(i)}). \end{aligned}$$

Therefore, it suffices to pre-compute $\{U(x_0^{(i)})\}_{i=1}^N$ to evaluate \hat{H} and its gradient efficiently.

The key quantity to control the accuracy is the variance of the approximation (27) which itself depends on (α, β) . The following provides a bound on the variance as a function of (α, β) . We defer its proof to Appendix D.1.

Proposition 8. Let $(\alpha, \beta) \in \mathbb{R} \times \mathbb{R}^q$ and let $x_0 \sim \Pi_\Omega$. Let $\sigma^2 = \sigma^2(\alpha, \beta)$ be the variance of the random variable $h(\alpha, \beta, x_0)$. Then

$$\sigma^2(\alpha, \beta) \leq 2|\alpha|^2 \|\beta\|_2^2 \mathbb{E}_{x'_0, x_0 \stackrel{iid}{\sim} \Pi_\Omega} [\|U(x_0) - U(x'_0)\|^2] + \|\beta\|_2^4 \mathbb{E}_{x'_0, x_0 \stackrel{iid}{\sim} \Pi_\Omega} [\|U(x_0)\|_2^2 \|U(x_0) - U(x'_0)\|_2^2].$$

Therefore, the approximation will tend to be accurate for small values of α and $\|\beta\|_2$. This can be used as a way to increase the number of samples in (27) depending on the magnitude of the approximate solution found.

In QMC the approximation in (27) is constructed from points $x_0^{(1)}, \dots, x_0^{(N)}$ that belong to a low-discrepancy sequence. The approximations converge faster than those obtained with MC, and significant improvements can be obtained in some cases; in others, the performance might be comparable to that of QMC [45]. There are standard constructions for low-discrepancy sequences in the hypercube, e.g., Halton, Sobol or Faure sequences. These can be used in the approximation (27) by scaling Ω and extending h by zero outside Ω . However, this leads to a possibly discontinuous integrand, which may substantially impact the performance of QMC. For this reason, the strategy introduced for the cubature formulae can be appropriate, in which we leverage Theorem 3 to represent H as an integral over the unit ℓ^2 -ball. In this case, we can use a classical one-dimensional sequence over the interval $[0, 1]$ to integrate \bar{h}° on r and then use the low-discrepancy sequence proposed in [46] to approximate the values of \bar{h}° .

5.2. Stochastic approximation methods

One of the drawbacks of (27) is that, similarly to cubature formulae, we may need a very large number of samples to approximate the objective function accurately. An alternative approach is to use *stochastic approximation* (SA) methods [44]. A popular SA method

is stochastic gradient descent (SGD) which uses the gradient of the integrand on a random sample at each iteration. Let $\{\mathbf{x}_0^{(i)}\}_{i \in \mathbb{N}}$ be an i.i.d. sequence of Π_Ω random vectors. Given an initial iterate (η_0, ξ_0) and a sequence of steps $\{s_n\}_{n \in \mathbb{N}}$ we define the iterates

$$\begin{aligned}\eta_{n+1} &= \eta_n - s_n(-1 + (\eta_n + \langle \mathbf{A}' \xi_n, \mathbf{U}(\mathbf{x}_0^{(n)}) \rangle_2)_+), \\ \xi_{n+1} &= \xi_n - s_n(\nabla L^\star(\lambda \xi_n, \mathbf{y}) + (\eta_n + \langle \mathbf{A}' \xi_n, \mathbf{U}(\mathbf{x}_0^{(n)}) \rangle_2)_+),\end{aligned}$$

for $n \geq 0$. In contrast to SAA, the samples change at each iteration. Hence, the computational cost is dominated by the evaluation of \mathbf{U} at a single point. However, the cost-per-iteration decreases at the expense of a decreased convergence rate.

One important factor impacting the performance of SGD is the choice of steps. It is known that if H is *strongly convex* then the steps may be chosen so that $s_n \sim n^{-1}$. However, as we do not know if the objective is strongly convex beforehand, we consider steps such that $s_n \sim n^{-1/2}$ [44].

In general, we may use a batch to estimate the gradient of H . Define the *unbiased* gradient estimate

$$\mathbf{g}_n(\alpha, \beta) = \frac{1}{M} \sum_{m=1}^M \nabla h(\alpha, \beta, \mathbf{x}_0^{(n,m)}) \quad (28)$$

where $\mathbf{x}_0^{(n,1)}, \dots, \mathbf{x}_0^{(n,M)} \stackrel{iid}{\sim} \Pi_\Omega$. One of the factors that impact the performance of SGD is the variance of this estimate. We defer the proof of the following proposition to Appendix D.2.

Proposition 9. *We have the bound*

$$\begin{aligned}\mathbf{E}_{\mathbf{x}_0^{(n,1)}, \dots, \mathbf{x}_0^{(n,M)} \sim \Pi_\Omega} [\|\mathbf{g}_n(\alpha, \beta) - \nabla H(\alpha, \beta)\|_2^2] &\leq \frac{1}{M} (|\alpha| + 6\|\beta\|_2^2) \mathbf{E}_{\mathbf{x}'_0, \mathbf{x}_0 \sim \Pi_\Omega} [\|\mathbf{U}(\mathbf{x}'_0) - \mathbf{U}(\mathbf{x}_0)\|_2^2] \\ &\quad + \frac{2}{M} \|\beta\|_2^2 \mathbf{E}_{\mathbf{x}'_0, \mathbf{x}_0 \sim \Pi_\Omega} [\|\mathbf{U}(\mathbf{x}'_0)\|_2^2 \|\mathbf{U}(\mathbf{x}'_0) - \mathbf{U}(\mathbf{x}_0)\|_2^2].\end{aligned}$$

Once again, the variance approximation becomes accurate for small values of α and $\|\beta\|_2$. This can be used as a way to select the size of the batch in (28) depending on the size of η and $\mathbf{A}' \xi$.

5.3. Samplers

Both in SAA and SGD it is necessary to be able to sample from Π_Ω . When Ω is convex and Π_Ω admits a density with respect to the Lebesgue measure, we can use *hit-and-run* [47] to sample efficiently from Π_Ω . This method is particularly efficient when Π_Ω is the uniform probability measure $\text{UNIF}(\Omega)$. One of its advantages is that the distribution of $\mathbf{x}_0^{(n)}$ converges efficiently to $\text{UNIF}(\Omega)$ in total variation, namely, in $O(d^3)$ steps [48, Corollary 1.2]. Furthermore, each iteration is computationally inexpensive: to sample uniformly from the unit sphere it suffices to sample from $\mathbf{g} \sim N(\mathbf{0}, \mathbf{I}_d)$ and let $\omega = \mathbf{g}/\|\mathbf{g}\|_2$, and for the sets Ω that we consider, finding the intersection between a line and its boundary is efficient. Remark that we may get good approximations for the integral even though the distribution of the current iterate may still be far from the uniform distribution.

Another alternative when Ω is simple, or the system is low-dimensional, is to use rejection sampling. This is particularly useful when Ω is an ℓ_p -ball as in this case we may use either an ℓ_∞ -ball or an ℓ_2 -ball containing it to sample from the corresponding uniform measure.

6. Experiments

6.1. A linear oscillator

As an illustrative example we consider a damped harmonic oscillator

$$\ddot{x} + \gamma \dot{x} + \omega^2 x = 0$$

with $x(0) = 1$ and $\dot{x}(1) = 0$. The reduction to a first-order system yields

$$\begin{cases} \dot{\mathbf{x}}(t) = \mathbf{f}(t, \mathbf{x}(t)) \\ \mathbf{x}(0) = \mathbf{x}_0 \end{cases} \quad \text{for } \mathbf{F} = \begin{bmatrix} 0 & 1 \\ -\omega^2 & -\gamma \end{bmatrix} \quad \text{and} \quad \mathbf{x}_0 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}. \quad (29)$$

This system has the form (2) for $\Pi_{x_0} = \delta_{[0, 1]^*}$ and it is instructive to study the probability densities that are obtained by the method in this case. Using an atomic measure to model uncertainty allows us to study whether the optimal density ρ^* is able to concentrate around this point. Assume that the QoIs are

$$u_1(\mathbf{x}_0) = x_1(2, \mathbf{x}_0), \quad u_2(\mathbf{x}_0) = x_2(4, \mathbf{x}_0), \quad u_3(\mathbf{x}_0) = x_1(6, \mathbf{x}_0), \quad u_4(\mathbf{x}_0) = x_2(8, \mathbf{x}_0).$$

Then \mathbf{U} can be written explicitly in terms of \mathbf{x}_0 as

Table 1

Bias and feasibility for the optimal probability measure.

Ω	r	$\ x_0 - E_{x_0 \sim \mu^*}[x_0]\ _2$				$\ U(x_0) - E_{x_0 \sim \mu^*}[U(x_0)]\ _2$			
		λ				λ			
		10 ⁰	10 ⁻¹	10 ⁻²	10 ⁻³	10 ⁰	10 ⁻¹	10 ⁻²	10 ⁻³
\bar{B}_1	0.5	9.76E-01	8.55E-01	7.07E-01	6.22E-01	7.32E-01	6.01E-01	4.88E-01	4.47E-01
	1.0	9.18E-01	6.58E-01	4.00E-01	2.26E-01	6.68E-01	4.26E-01	2.25E-01	1.64E-01
	2.0	7.60E-01	3.34E-01	8.98E-02	4.40E-02	5.19E-01	1.90E-01	3.42E-02	2.02E-02
\bar{B}_2	0.5	9.67E-01	8.03E-01	6.61E-01	5.82E-01	7.24E-01	5.59E-01	4.26E-01	3.81E-01
	1.0	8.87E-01	5.61E-01	3.17E-01	1.82E-01	6.35E-01	3.67E-01	1.83E-01	1.32E-01
	2.0	6.91E-01	2.26E-01	4.60E-02	6.76E-03	4.49E-01	1.37E-01	4.82E-02	1.39E-02
\bar{B}_∞	0.5	9.56E-01	7.64E-01	6.19E-01	5.63E-01	7.09E-01	5.16E-01	3.84E-01	3.14E-01
	1.0	8.58E-01	4.89E-01	2.91E-01	1.15E-01	6.13E-01	2.95E-01	1.64E-01	8.82E-02
	2.0	6.34E-01	1.75E-01	2.02E-02	1.49E-02	4.04E-01	1.03E-01	9.87E-03	1.56E-02
\bar{B}_1^+	0.5	8.40E-01	7.61E-01	6.16E-01	5.56E-01	7.02E-01	6.16E-01	4.85E-01	4.03E-01
	1.0	6.92E-01	4.80E-01	2.99E-01	1.87E-01	6.38E-01	4.31E-01	2.42E-01	1.34E-01
	2.0	5.79E-01	3.02E-01	2.20E-01	7.57E-02	5.49E-01	2.55E-01	1.15E-01	6.66E-02
\bar{B}_2^+	0.5	8.04E-01	7.14E-01	5.86E-01	5.48E-01	6.95E-01	6.07E-01	5.03E-01	4.00E-01
	1.0	6.63E-01	4.40E-01	2.60E-01	1.84E-01	6.41E-01	4.28E-01	2.58E-01	1.82E-01
	2.0	6.92E-01	4.08E-01	2.12E-01	9.71E-02	5.94E-01	2.77E-01	1.32E-01	6.29E-02
\bar{B}_∞^+	0.5	7.79E-01	6.99E-01	5.83E-01	5.81E-01	6.94E-01	6.02E-01	4.88E-01	4.62E-01
	1.0	6.56E-01	4.40E-01	2.62E-01	1.85E-01	6.50E-01	4.42E-01	2.32E-01	1.58E-01
	2.0	8.26E-01	4.59E-01	2.27E-01	1.14E-01	6.35E-01	3.15E-01	1.46E-01	7.23E-02

$$U(x_0) = \begin{bmatrix} \langle e_1, e^F x_0 \rangle_2 \\ \langle e_2, e^{2F} x_0 \rangle_2 \\ \langle e_1, e^{4F} x_0 \rangle_2 \\ \langle e_2, e^{6F} x_0 \rangle_2 \end{bmatrix} = \begin{bmatrix} e_1^* e^F \\ e_2^* e^{2F} \\ e_1^* e^{4F} \\ e_2^* e^{6F} \end{bmatrix} x_0 =: Ux_0. \quad (30)$$

Fig. 2a shows how x and \dot{x} evolve in time, whereas Fig. 2b shows the trajectory of the system in state space. Fig. 2c shows the QoI being measured. The QoI (30) can be computed efficiently and with high accuracy, i.e., the matrix F is non-defective and its exponential can be computed from its eigendecomposition, allowing us to focus on the density obtained with the method. To compare the effect that Ω has on the results, we consider 6 possible convex bodies: the closed ℓ_1 -ball \bar{B}_1 , the closed ℓ_2 -ball \bar{B}_2 , the closed ℓ_∞ -ball \bar{B}_∞ , and their intersections with the positive orthant, which we denote as \bar{B}_1^+ , \bar{B}_2^+ and \bar{B}_∞^+ . In all cases we assume that μ_Ω is the uniform measure. These choices will help us evaluate the impact of incorporating *a priori* information on Ω , i.e., the fact that $x_0 \geq 0$. To approximate H we use Monte Carlo and we draw 10,000 samples from μ_Ω in each case using rejection sampling.

In Fig. 3 we show the densities obtained for \bar{B}_2 . As U is a linear map, in all cases the density is the positive part of an affine function. For $r = 1/2$ we have that $x_0 \notin \Omega$ and the problem becomes infeasible as $\lambda \rightarrow 0$. In Figs. 3a-3d we see that the density concentrates in the boundary as $\lambda \rightarrow 0$. In addition, there seems to be a small bias towards negative values of \dot{x} . This is a consequence of the particular structure of the matrix U . For $r = 1$ we have that $x_0 \in \text{bd}(\Omega)$. As shown in Proposition 6, this implies that the problem becomes infeasible as $\lambda \rightarrow 0$. Figs. 3e-3h show the optimal density concentrating near the boundary. Observe that, as $\lambda \rightarrow 0$, the expected value of Π^* approaches the true x_0 . Finally, for $r = 2.0$ we have $x_0 \in \text{int}(\Omega)$. As Figs. 3i-3l show, when $\lambda \rightarrow 0$ the density somewhat concentrates around x_0 and the expected value quickly approaches the true x_0 . However, as U is linear, the set of possible densities is not flexible enough to truly concentrate around the true x_0 , generating instead a large dispersion around it.

In Fig. 4 we show the densities obtained for $r = 1$ and all of the 6 choices of Ω . Remark that in all cases we have that $x_0 \in \text{int}(\Omega)$. In general, the behavior as $\lambda \rightarrow 0$ is the same observed for the ℓ_2 -ball, and the density concentrates near the boundary and near x_0 . However, the behavior of the optimal density behaves quite differently depending on the choice of Ω . In fact, as $\lambda \rightarrow 0$ the dispersion on the marginal for \dot{x} depends almost exclusively on the choice of Ω , yielding the least dispersion for \bar{B}_1 , and the largest for \bar{B}_∞ . By considering the closed ℓ_p -balls intersected with the positive orthant, we are adding *a priori* information about the values of $x(0)$ and $\dot{x}(0)$. However, this leads to a bias in the expected value of Π^* . The magnitude of this bias and its behavior as $\lambda \rightarrow 0$ depends on the structure of the boundary of Ω at x_0 .

In Table 1 we report the ℓ_2 -bias when estimating x_0 by the expected value of Π^* and the feasibility error between the true expected value $U(x_0)$ and that obtained using Π^* .

This example also allows us to compare the densities obtained by solving the finite-dimensional approximations discussed in Sections 2.2 and 3.1. For this comparison we consider $\Omega = \bar{B}_\infty$ and $\lambda = 10^{-1}$, and we use the prior induced by the Haar wavelet basis on $[-1, 1]^2$. A crucial property of the Haar wavelet basis is that it is piecewise constant, enabling the efficient characterization of the set Θ_{V_n} . The results obtained from the finite-dimensional approximations up to a level $L \in \{1, 2, 3\}$ can be seen in Fig. 5; recall that there are 4^{L+1} basis functions for each choice of L . As the decomposition level increases, the finite-dimensional approximation begins to resemble the solution shown in Fig. 4i. In Fig. 5d we see that the L^2_μ -distance decreases, as expected.

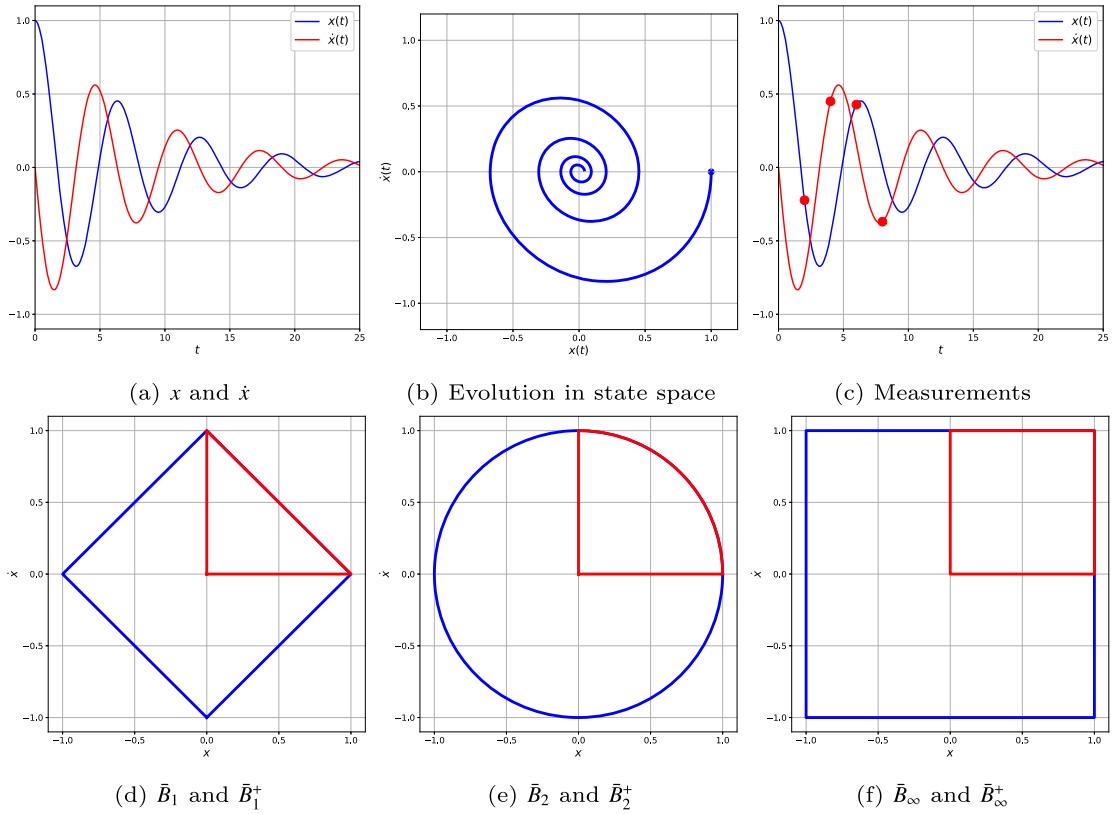


Fig. 2. (a) Evolution of the state variables x and \dot{x} on the time interval $[0, 20]$ for the damped oscillator (29) for $\gamma = 1/4$ and $\omega = 1$. (b) Trajectory in phase space for $t \in [0, 20]$. (c) Observed values of the state variables. (d, e, f) Diagram of \bar{B}_1 , the \bar{B}_2 and \bar{B}_∞ in blue, and \bar{B}_1^+ , \bar{B}_2^+ and \bar{B}_∞^+ in red.

Therefore, the finite-dimensional approximations do converge to the solution to the infinite-dimensional problem. This validates our interpretation of the problem as the limit of the finite-dimensional models.

6.2. Failure on a linear flow system

Another illustrative example is to consider a linear flow system on a graph (Fig. 6). At the initial time the first node always has concentration equal to 1 whereas all other nodes always have concentration equal to zero. The *normal state* of the system occurs when the flow goes through node 2 and ends up at node 3, i.e., the capacity of the edges $2 \rightarrow 4$ and $4 \rightarrow 2$ are zero (Fig. 6a). The *failure state* occurs when the capacity of the edges $2 \rightarrow 4$ and $4 \rightarrow 2$ are non-zero and, as a consequence, there is a leak from node 2 to node 4 (Fig. 6b). The goal of this problem is to quantify the uncertainty on both on the initial condition and on the parameters $c_{2 \rightarrow 4}$ and $c_{4 \rightarrow 2}$ from measurements of the concentration at nodes 1 and 3 at times $\{1.0, 2.0, \dots, 9.0, 10.0\}$.

We consider three scenarios of failure. In *Scenario A* the trajectories are randomly generated according to $c_{2 \rightarrow 4} \sim \Pi_{2 \rightarrow 4}$ with $\Pi_{2 \rightarrow 4} = \text{BETA}(30, 10)$ and $c_{4 \rightarrow 2} \sim \Pi_{4 \rightarrow 2}$ for $\Pi_{2 \rightarrow 4} = \text{BETA}(15/0.85, 10)$; in this case $E[c_{2 \rightarrow 4}] = 0.75$ and $E[c_{4 \rightarrow 2}] = 0.15$ (Fig. 6c). The sample trajectories and the average measurements are shown in Fig. 6d. In *Scenario B* the trajectories are generated from a mixture with probability 1/2 between the normal state, i.e., $c_{2 \rightarrow 4} = c_{4 \rightarrow 2} = 0$, and the trajectories in Scenario B (Fig. 6e). Finally, in *Scenario C* we consider that the system operates normally for $t \leq 2.0$, but then fails for $t > 2.0$ where the parameters distribute as in Scenario A (Fig. 6f). In each case, we use the average observations over 100 samples.

To use the method we represent the model as the system of ODEs

$$\begin{aligned}
 \dot{v}_1(t) &= -c_{1 \rightarrow 2}v_1(t) \\
 \dot{v}_2(t) &= -(c_{2 \rightarrow 3} + c_{2 \rightarrow 4}(t))v_2(t) + c_{1 \rightarrow 2}v_1(t) + c_{4 \rightarrow 2}(t)v_4(t) \\
 \dot{v}_3(t) &= c_{2 \rightarrow 3}v_2(t) \\
 \dot{v}_4(t) &= -c_{4 \rightarrow 2}(t)v_4(t) + c_{2 \rightarrow 4}(t)v_2(t) \\
 \dot{c}_{2 \rightarrow 4}(t) &= 0 \\
 \dot{c}_{4 \rightarrow 2}(t) &= 0.
 \end{aligned}$$

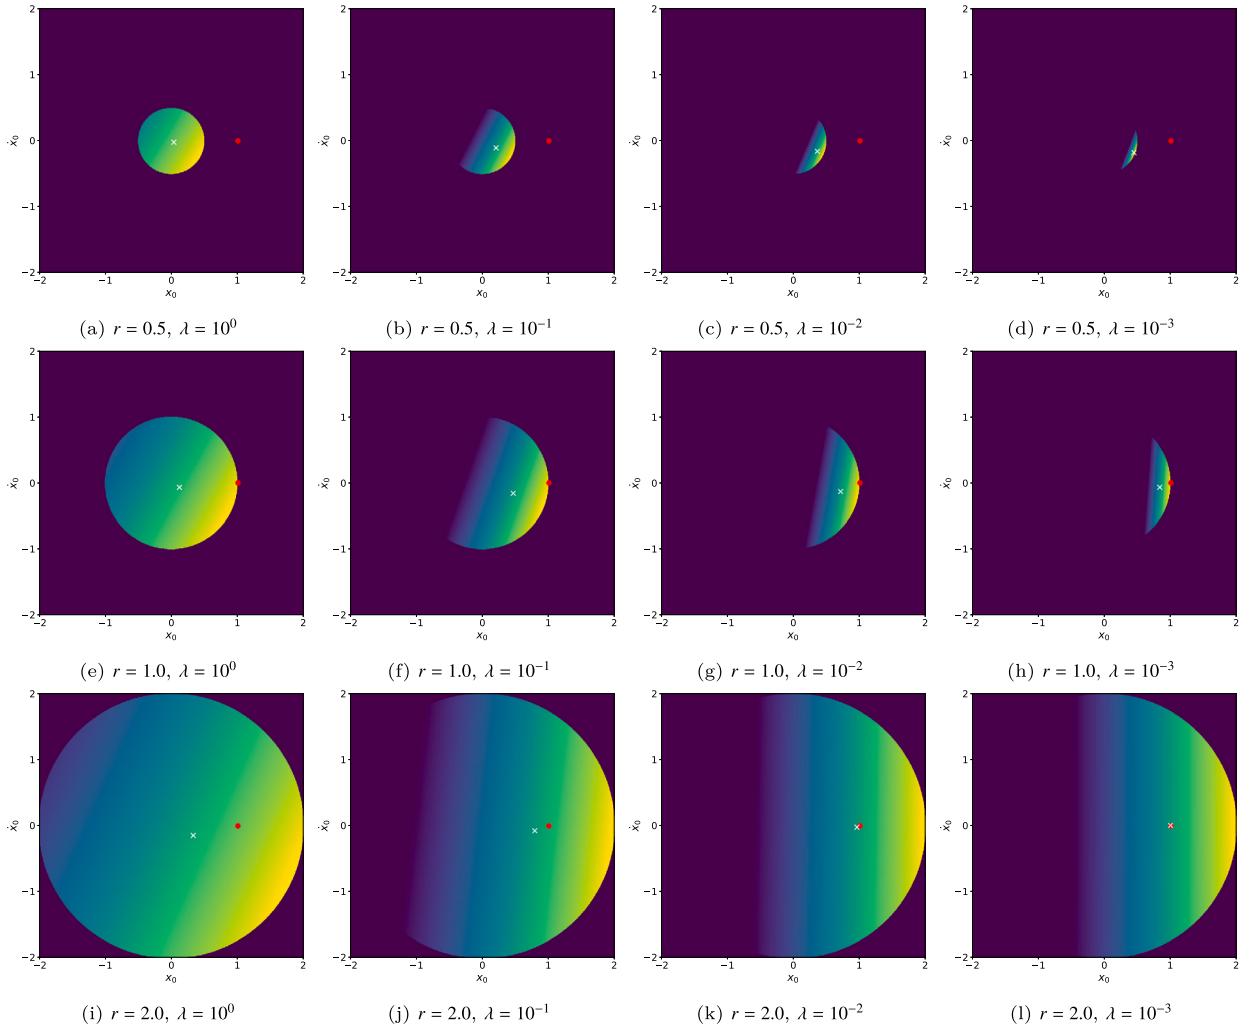


Fig. 3. Optimal density ρ^* for \bar{B}_2 for radii $r \in \{0.5, 1.0, 2.0\}$ and $\lambda \in \{10^0, 10^{-1}, 10^{-2}, 10^{-3}\}$. In each case, the values of the density have been normalized to the interval $[0, 1]$, the true x_0 is represented as a red dot, and the expected value for ρ^* is shown as a white dot.

If we let

$$\boldsymbol{\nu} = \begin{bmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \\ \nu_4 \end{bmatrix}, \quad \boldsymbol{\vartheta} = \begin{bmatrix} c_{2 \rightarrow 4} \\ c_{4 \rightarrow 2} \end{bmatrix} \quad \text{and} \quad \mathbf{x} = \begin{bmatrix} \boldsymbol{\nu} \\ \boldsymbol{\vartheta} \end{bmatrix}$$

then we have the non-linear system

$$\begin{cases} \dot{\mathbf{x}}(t) = \mathbf{C}(\mathbf{x}(t))\mathbf{x}(t) \\ \mathbf{x}(0) = \mathbf{x}_0 \end{cases} \quad (31)$$

where

$$\mathbf{C}(\mathbf{x}) = \begin{bmatrix} -c_{1 \rightarrow 2} & 0 & 0 & 0 & 0 & 0 \\ c_{1 \rightarrow 2} & -(c_{2 \rightarrow 3} + x_5) & 0 & x_6 & 0 & 0 \\ 0 & c_{2 \rightarrow 3} & 0 & 0 & 0 & 0 \\ 0 & x_5 & 0 & -x_6 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

We consider the QoIs $\mathbf{U} : \mathbb{R}^6 \rightarrow \mathbb{R}^{20}$ given by

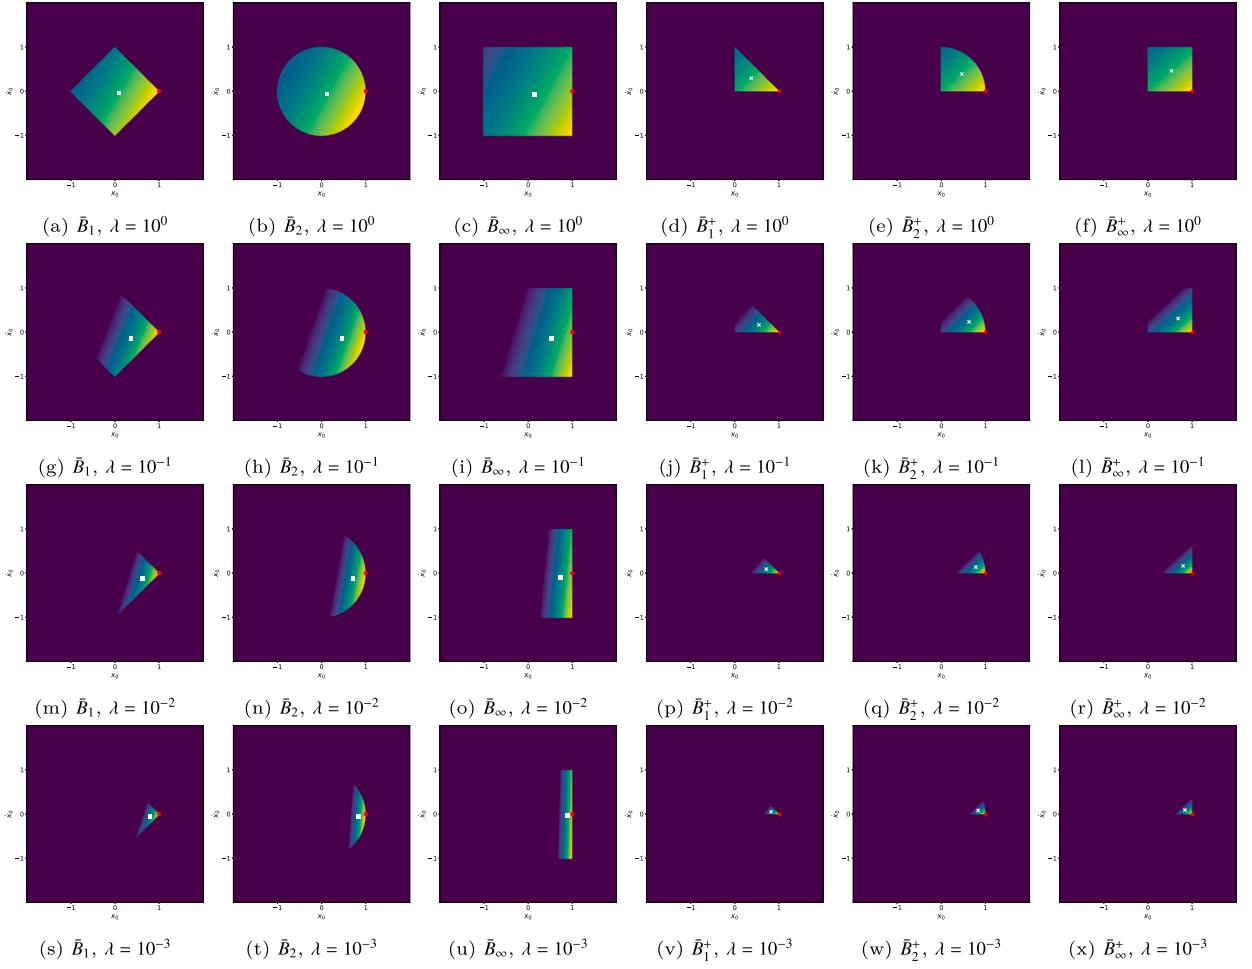


Fig. 4. Optimal density ρ^* for $r = 1.0$ and $\lambda \in \{1^0, 10^{-1}, 10^{-2}, 10^{-3}\}$. In each case, the values of the density have been normalized to the interval $[0, 1]$, the true \mathbf{x}_0 is represented as a red dot, and the expected value for ρ^* is shown as a white dot.

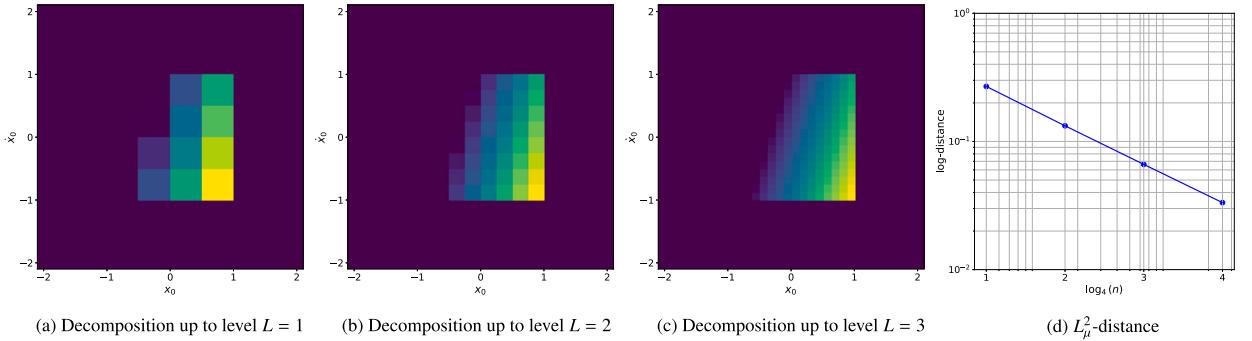


Fig. 5. Solutions to the finite-dimensional problems using the Haar wavelet basis on \bar{B}_∞ . For each level L of the decomposition the corresponding subspace V_n has dimension $n = 4^{L+1}$. The L_μ^2 -distance between the optimal density ρ^* obtained by solving the infinite-dimensional problem and the finite-dimensional problems decreases as expected.

$$U_1(\mathbf{x}_0) = v_1(1, \mathbf{x}_0), \dots, U_{10}(\mathbf{x}_0) = v_1(10, \mathbf{x}_0), U_{11}(\mathbf{x}_0) = v_3(1, \mathbf{x}_0), \dots, U_{20}(\mathbf{x}_0) = v_3(10, \mathbf{x}_0)$$

and the measure $\Pi_{\mathbf{x}_0}$ on \mathbb{R}^6 for \mathbf{x}_0 is

$$\Pi_{\mathbf{x}_0} = \delta_1 \otimes \delta_0 \otimes \delta_0 \otimes \delta_0 \otimes \pi_{2 \rightarrow 4} \otimes \pi_{4 \rightarrow 2}.$$

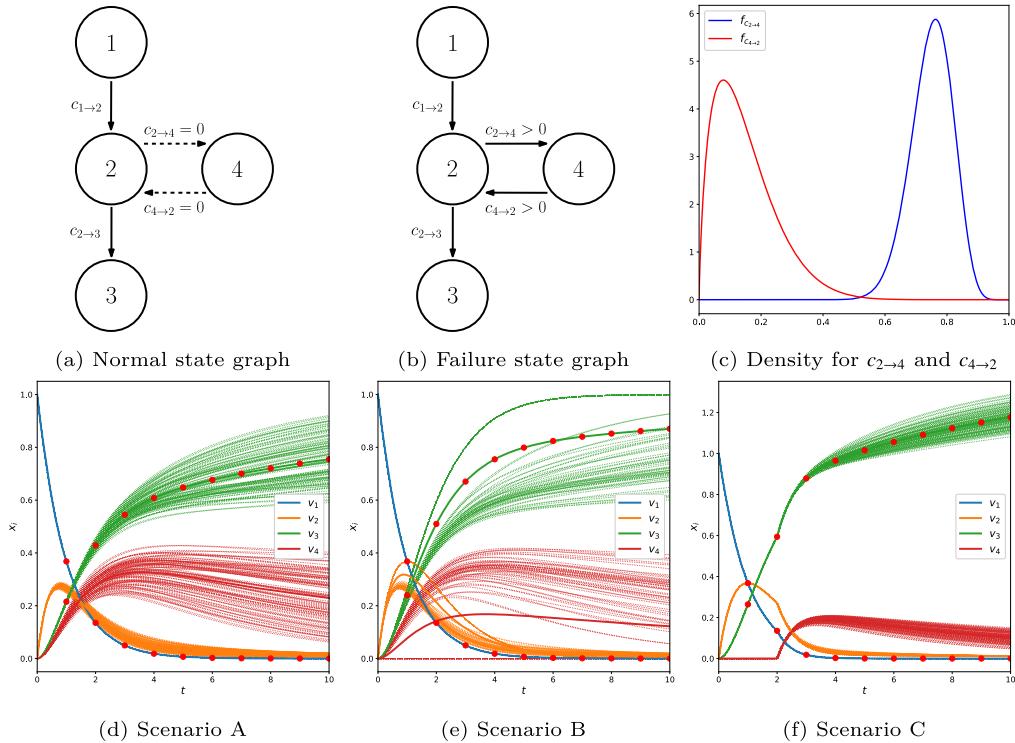


Fig. 6. (a, b) Graph representing the normal and failure states of the system. (c) Probability densities for $c_{2 \rightarrow 4}$ and $c_{4 \rightarrow 2}$. (d, e, f) Trajectories generated from 100 samples of $c_{2 \rightarrow 4}$ and $c_{4 \rightarrow 2}$ under each of the failure scenarios. In each case, the solid line represents the average trajectory whereas the red dots represent the average measurements.

For the numerical experiments, we use $c_{1 \rightarrow 2} = c_{2 \rightarrow 3} = 1$. As the solutions to the system diverge exponentially if $c_{2 \rightarrow 4} < 0$ or $c_{4 \rightarrow 2} < 0$ it is reasonable to impose that the support of the measure on the last 2 variables is contained in the positive orthant. Consequently, to perform the experiments we choose Ω as either $r\bar{B}_1^+$, $r\bar{B}_2^+$ or $r\bar{B}_\infty^+$, or as Cartesian products of the form $r\bar{B}_{p_1} \times r\bar{B}_{p_2}^+$ for $p_1, p_2 \in \{1, 2, \infty\}$ and some radius $r > 0$. To avoid the boundary effects shown in Section 6.1 we choose $r = 5$. To approximate \bar{H} we use Monte Carlo and we sample 10,000 from μ_Ω in each case using rejection sampling. In all cases we let μ_Ω be the uniform measure on Ω and we let $\lambda = 10^{-3}$. Finally, we solve numerically (31) using SciPy's implementation of the RK45 method with absolute tolerance 10^{-4} and relative tolerance 10^{-6} .

If Fig. 7 we show numerical approximations for the marginals for Π^* for the selected choices $5\bar{B}_1^+$, $5\bar{B}_2 \times 5\bar{B}_1^+$ and $5\bar{B}_1^+ \times 5\bar{B}_2^+$ for Ω . To approximate the marginals for variables taking values in $[0, 5]$ we divided the interval in 10 bins, whereas for variables taking values in $[-5, 5]$ we divided the interval in 20 bins. The probability of each bin under Π^* was approximated with 20,000 Monte Carlo samples independent from those used to approximate H .

The marginals obtained for $5\bar{B}_1^+$ for the components of v_0 concentrate near zero, except the marginal for $v_{0,1}$ (rows 1 to 3 of Fig. 7). However, the marginals cannot concentrate sufficiently near zero. As a consequence the expected values of $v_{0,2}$ and $v_{0,3}$ tend to be too large. This may also be the reason why the marginals for $v_{0,4}$ assign large probability to large values of $v_{0,4}$. In contrast, the marginals for $c_{2 \rightarrow 4}$ and $c_{4 \rightarrow 2}$ capture the behavior of these variables reasonably well, as the marginal for $c_{4 \rightarrow 2}$ is more concentrated toward the origin than the marginal for $c_{2 \rightarrow 4}$. This behavior is consistent for all scenarios. Finally, in Figs. 8a-8c we see the trajectory associated to the expected value of v_0 , $c_{2 \rightarrow 4}$ and $c_{4 \rightarrow 2}$. Although they capture reasonably well the measurements of v_1 in time, they tend to overestimate the average measurements for v_3 . This is itself a consequence of overestimating $v_{0,4}$.

Observe that the marginals obtained for $5\bar{B}_2 \times 5\bar{B}_1^+$ for the components of v_0 spread over the interval $[-5, 5]$ but concentrate near the origin (rows 4 to 6 of Fig. 7). As a consequence, the probability of observing a negative value is large. However, the expected values for $v_{0,2}$ and $v_{0,3}$ are closer to the initial condition of the system for Scenarios A and B; for Scenario C the expected value for $v_{0,3}$ is negative and that for $v_{0,4}$ is much larger. This suggests the dynamics of the system constrain these two variables. In all cases, the marginals for $c_{2 \rightarrow 4}$ and $c_{4 \rightarrow 2}$ behave in a similar manner, suggesting that in this setup the method cannot differentiate between the effect of these two variables. Figs. 8d-8f show the trajectory associated to the expected values of x_0 , $c_{2 \rightarrow 4}$ and $c_{4 \rightarrow 2}$. Choosing Ω as a Cartesian product allows for more flexibility, and as a consequence, the trajectories seem to be a good fit to the observations in all scenarios. However, note that the initial condition may have negative entries. This effect can be mitigated by choosing a suitable set C in (13).

Finally, the marginals obtained for $5\bar{B}_1^+ \times 5\bar{B}_2^+$ for the components of x_0 behave similarly to those obtained for $5\bar{B}_1^+$ (rows 7 to 9 in Fig. 7). Remark that for this choice the marginals for $v_{0,1}$ and $v_{0,4}$ behave similarly and their expected values are also similar. As for $5\bar{B}_1^+$ we also see that the densities for $c_{2 \rightarrow 4}$ and $c_{4 \rightarrow 2}$ then to capture the overall relation between $c_{2 \rightarrow 4}$ and $c_{4 \rightarrow 2}$. In Figs. 8g-8i the

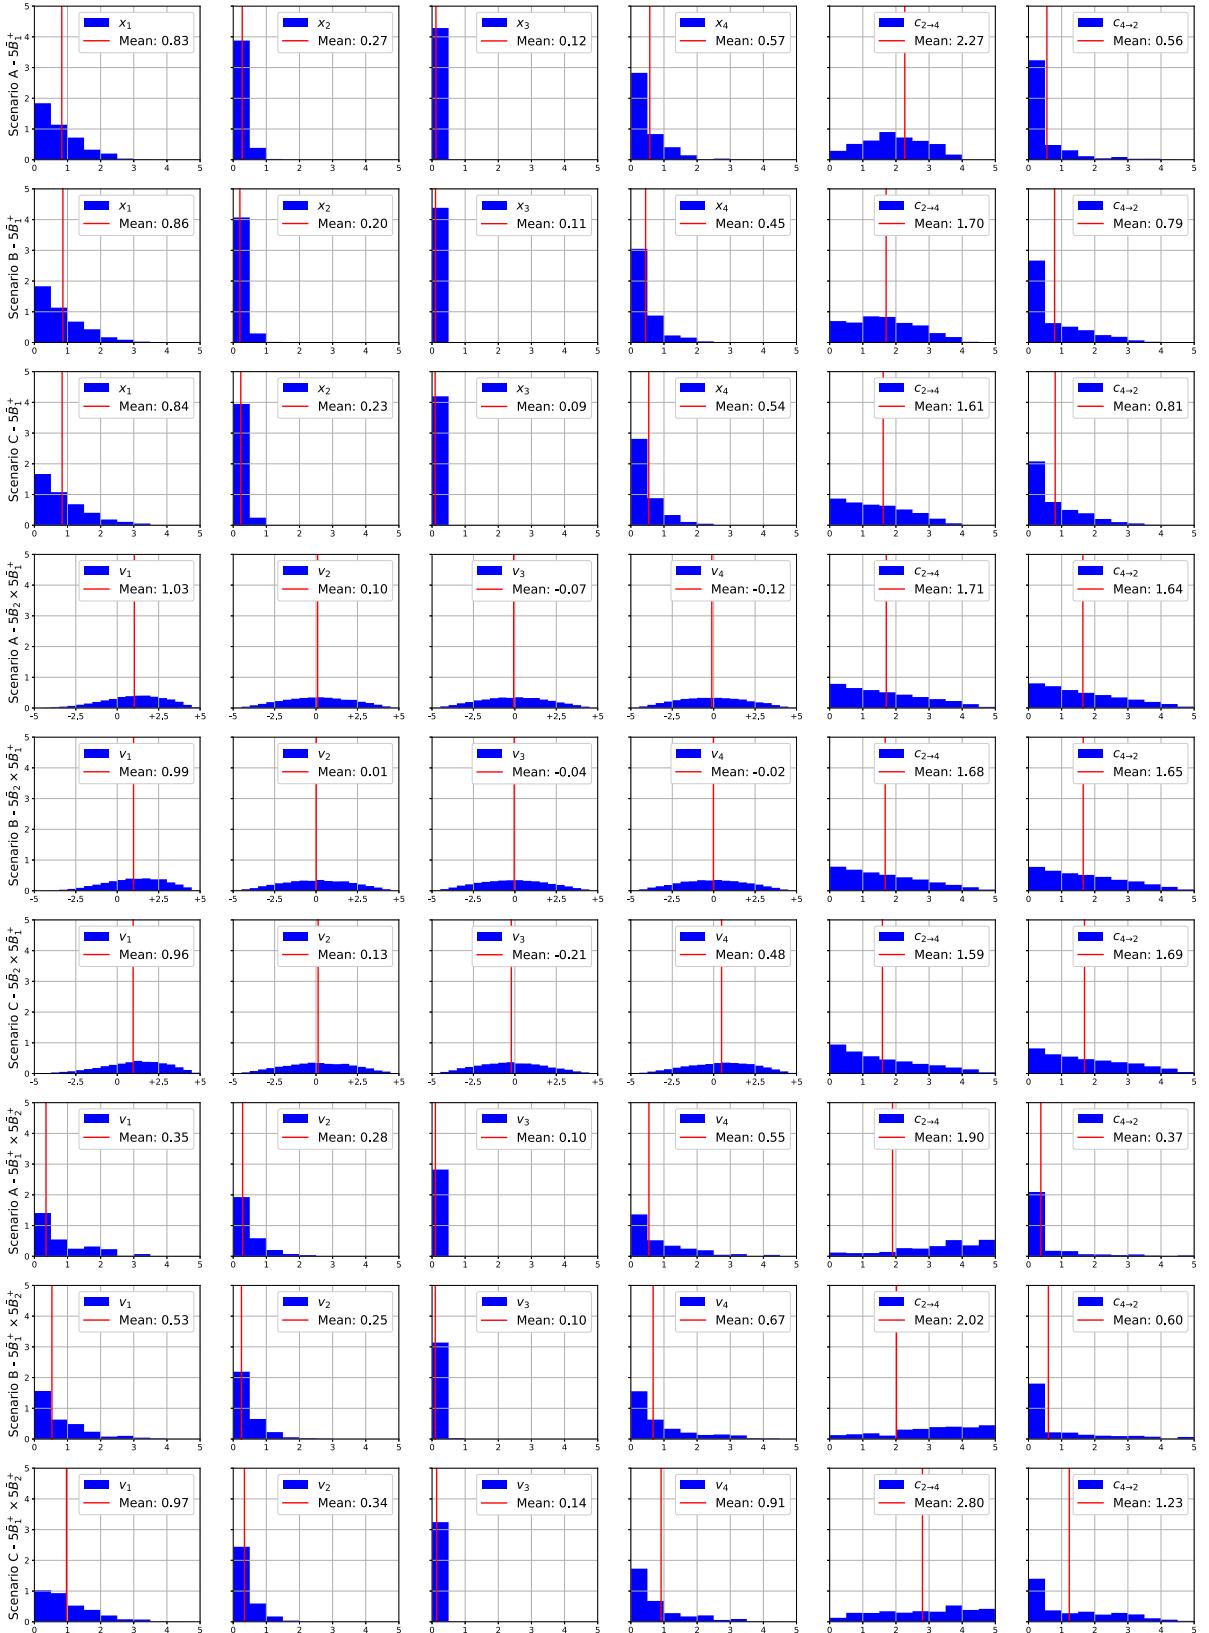
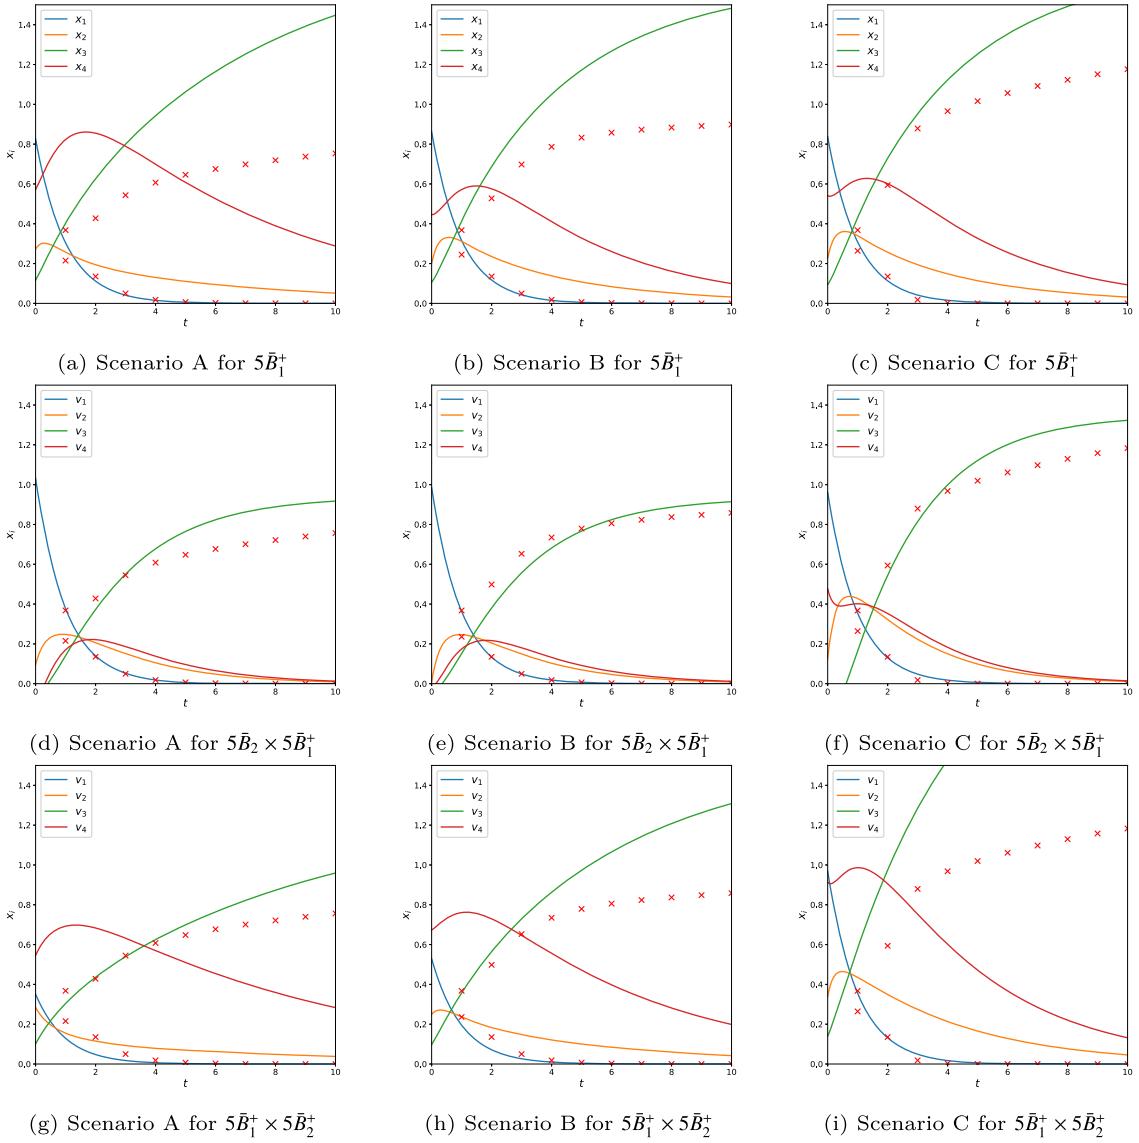


Fig. 7. Marginals for Π^* for selected choices of Ω under each of the failure scenarios. Rows 1 to 3 show the results for $5\bar{B}_1^+$, rows 4 to 6 for $5\bar{B}_2 \times 5\bar{B}_1^+$ and rows 7 to 9 for $5\bar{B}_1^+ \times 5\bar{B}_2^+$. The vertical line indicates the expected value.

Fig. 8. Trajectories for the expected value $E_{x_0 \sim \Pi^*} [x_0]$.

trajectories for the expected value show similarities with those obtained for $5\bar{B}_1^+$. However, for Scenario C the differences between this trajectory and the average observations for v_3 are quite large. This is due to the fact that an overestimate of $v_{0,2}$ has a strong impact on the limiting value of v_3 .

6.3. Chemical reaction network

To illustrate the application of our method to a high-dimensional, non-linear system, we consider a chemical reaction network. These can be described by the reaction rate equations, which are a non-linear system of ODEs that model the time-evolution of the concentrations of the chemical species in the system [49]. If the system has n species and we let $c \in \mathbb{R}^n = c(t)$ denote the concentrations, then the reaction rate equations become

$$\dot{c}(t) = \sum_{i=1}^p k_i h_i(c(t)) \mathbf{r}_i \quad (32)$$

where p is the number of chemical reactions in the system, $k_1, \dots, k_p > 0$ are the reaction *rate constants*, $h_1, \dots, h_p : \mathbb{R}^n \rightarrow \mathbb{R}$ are the *propensity functions* and $\mathbf{r}_1, \dots, \mathbf{r}_p \in \mathbb{R}^n$ are the *stoichiometry vectors*. In this case, we shall assume the initial concentrations c_0 are fixed. However, the rate constants will be random. We assume $k_i := \delta_i k_i^*$ where $k_i^* > 0$ is a known reference value for the rate constant of the i -th reaction and $\delta_1, \dots, \delta_p$ are i.i.d. $\text{LOGNORMAL}(-\sigma^2/2, \sigma)$ for $\sigma = \log(10)/2$. This ensures the rates vary roughly across 2 orders of magnitude and that the expected value of k_i is close to its reference value k_i^* .

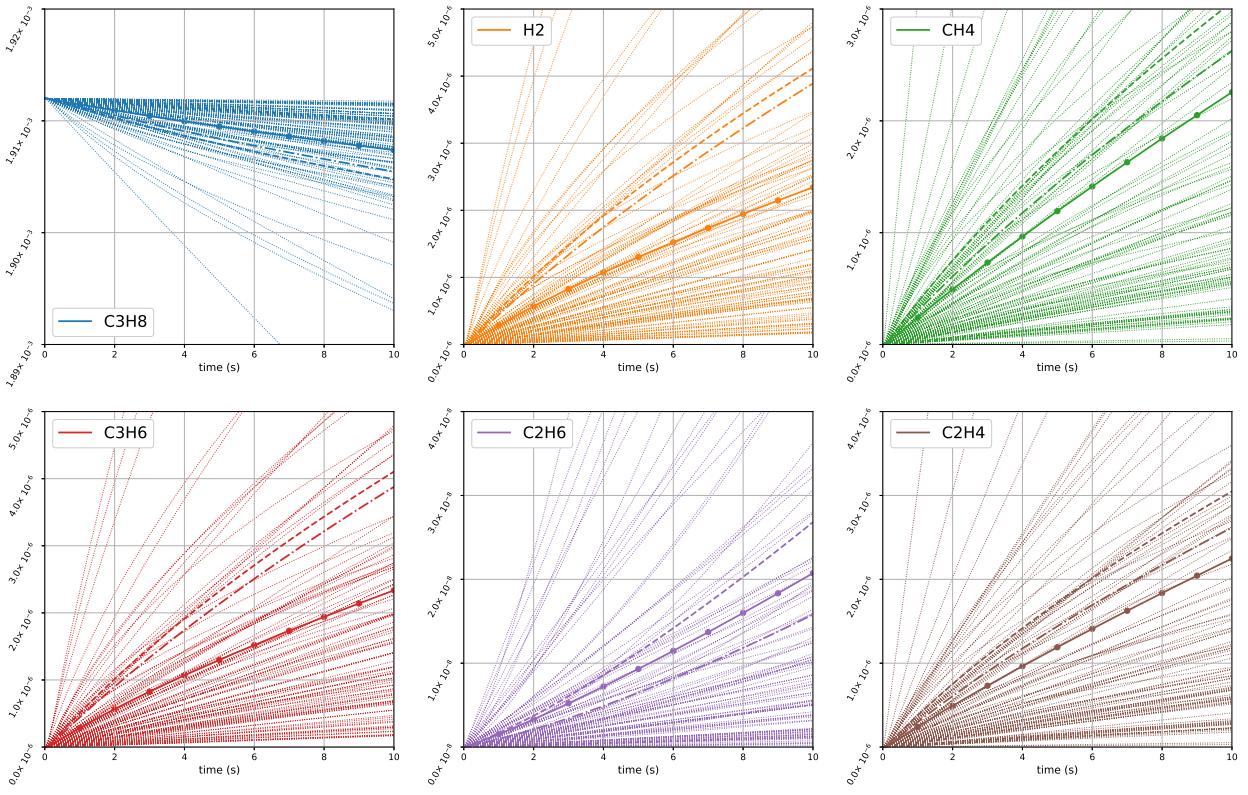


Fig. 9. Concentrations for the observed species C_3H_8 , H_2 , CH_4 , C_3H_6 , C_2H_6 and C_2H_4 . The dotted lines represent sample trajectories. The solid line represents the average trajectory, while the circles represent the average value of the observations. The dashed line represents the trajectory associated to the reference values k^* . The dashed-dotted line represents the trajectory associated to the average rates $\mathbf{E}_{k \sim \mu_0}[k]$.

We choose $\Omega = [10^{-2}k_1^*, 10^{+2}k_1^*] \times \dots \times [10^{-2}k_p^*, 10^{+2}k_p^*]$ and μ_Ω as the joint density of the random variables $10^{\alpha_1}k_1^*, \dots, 10^{\alpha_p}k_p^*$ for $\alpha_1, \dots, \alpha_p \stackrel{iid}{\sim} \text{UNIF}([-2, 2])$. To perform the simulations, we consider a medium-sized reaction system modeling propane pyrolysis with 38 chemical species and 98 reactions [50, Section IV.B]. We simulate the system for $t \in [0, 10]$ and observe every $\Delta t = 1$ starting at $t = 1$ the concentrations of the 6 most chemically relevant species in the system: C_3H_8 , H_2 , CH_4 , C_3H_6 , C_2H_6 and C_2H_4 (Fig. 9). We sample 100 trajectories to compute the average observation. The trajectories are found by solving (32) numerically using SciPy's implementation of the Radau method with absolute tolerance 10^{-8} and relative tolerance 10^{-10} . To approximate the objective function, we use 10,000 samples from μ_Ω .

In Fig. 10 we show the marginals obtained for 6 chemically relevant reactions $\text{C}_3\text{H}_8 \rightarrow \text{CH}_3 + \text{C}_2\text{H}_5$, $\text{CH}_3 + \text{H}_2 \rightarrow \text{H} + \text{CH}_4$, $\text{C}_3\text{H}_5 + \text{H}_2 \rightarrow \text{H} + \text{C}_3\text{H}_6$, $\text{H} + \text{C}_3\text{H}_6 \rightarrow \text{C}_3\text{H}_5 + \text{H}_2$, $\text{C}_2\text{H}_5 + \text{H}_2 \rightarrow \text{H} + \text{C}_2\text{H}_6$ and $\text{CH}_3 + \text{C}_2\text{H}_6 \rightarrow \text{C}_2\text{H}_5 + \text{CH}_4$ for $\lambda = 10^{-2}$. As we can see, the marginal remains almost flat over the interval. This may be a consequence of the fact that the system is ill-conditioned, that the variations of \mathbf{U} are small, and that the amount of measurements is not sufficient to accurately characterize the uncertainty. Furthermore, there is a bias in the estimate. In Fig. 11 we see the sample trajectories from the optimal density Π^* .

7. Discussion

Our experiments show that the method we proposed to compute the optimal density can be implemented efficiently. For typical choices of Ω for low-dimensional problems, rejection sampling and stochastic approximation methods are sufficient to obtain good results. For high-dimensional problems, stochastic gradient descent may perform well. In this case, the main cost is evaluating the map \mathbf{U} . Furthermore, the closed-form expression for the optimal density allows us to compute statistics that may be relevant when quantifying the uncertainty in the quantities of interest. In particular, this allows us to compute the marginal distributions of these quantities.

We have seen that the choice of Ω has a strong impact on the results. Although choosing a “small” Ω could have computational advantages when approximating H it may render the problem ill-conditioned as was shown in Section 6.1. When the density accumulates near the boundary of Ω its geometry will strongly influence the structure of, e.g., the marginals. The results in Section 6.2 show that even if the densities concentrate near the boundary, the expected value may be biased. Furthermore, the geometry of the boundary of Ω will determine if the density is such that the marginals concentrate for all or only some of the state variables.

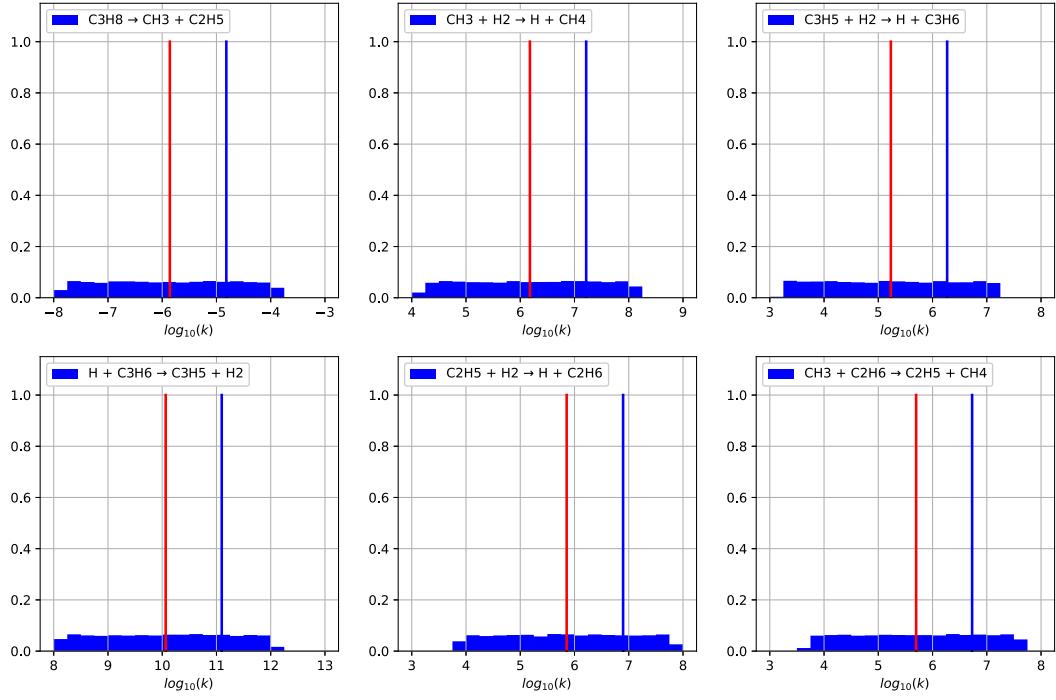


Fig. 10. Histograms for the optimal marginal densities for $\lambda = 10^{-2}$ and the selected reactions $\text{C}_3\text{H}_8 \rightarrow \text{CH}_3 + \text{C}_2\text{H}_5$, $\text{CH}_3 + \text{H}_2 \rightarrow \text{H} + \text{CH}_4$, $\text{C}_3\text{H}_5 + \text{H}_2 \rightarrow \text{H} + \text{C}_3\text{H}_6$, $\text{H} + \text{C}_3\text{H}_6 \rightarrow \text{C}_3\text{H}_5 + \text{H}_2$, $\text{C}_2\text{H}_5 + \text{H}_2 \rightarrow \text{H} + \text{C}_2\text{H}_6$ and $\text{CH}_3 + \text{C}_2\text{H}_6 \rightarrow \text{C}_2\text{H}_5 + \text{CH}_4$. The densities are in the logarithm of the rate. As we can see, the density remains almost uniform over the interval of interest.

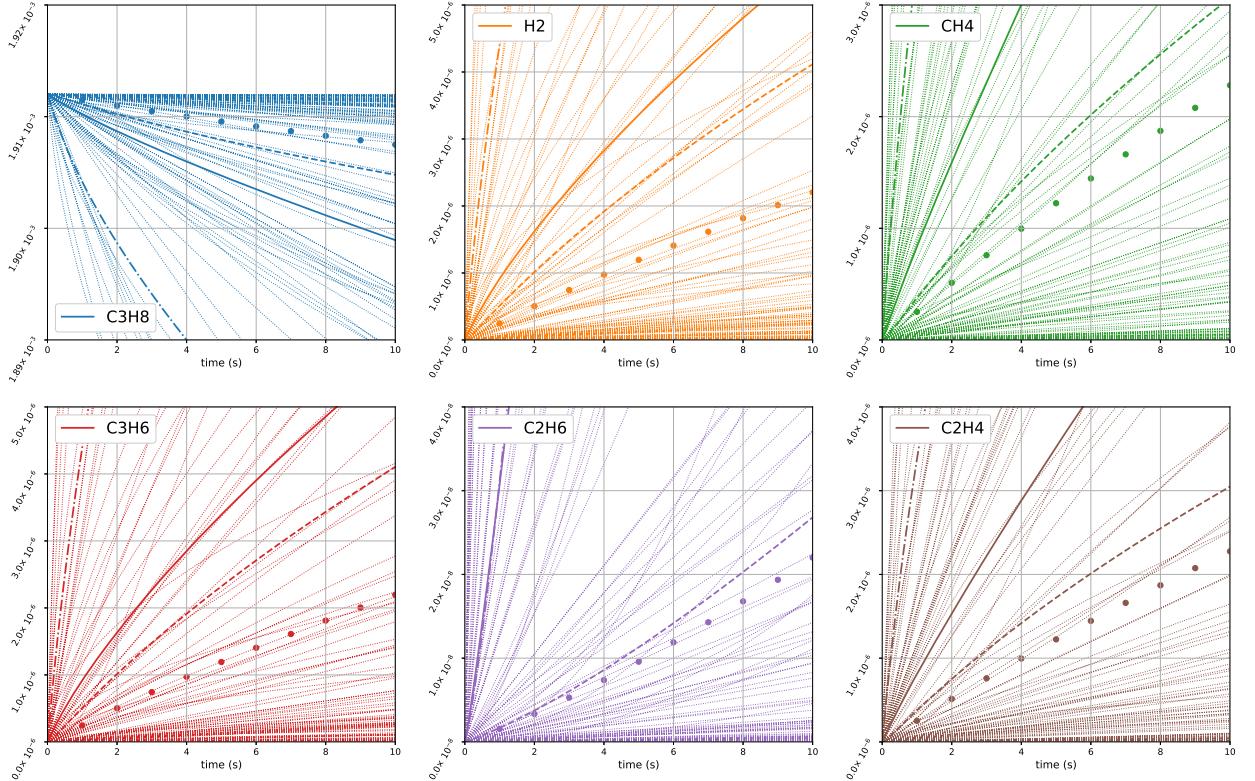


Fig. 11. Concentrations for the observed species C_3H_8 , H_2 , CH_4 , C_3H_6 , C_2H_6 and C_2H_4 for the optimal density for $\lambda = 10^{-2}$. The dotted lines represent sample trajectories from the optimal density. The solid line represents the average trajectory, while the circles represent the average value of the observations. The dashed line represents the trajectory associated to the reference values k^* . The dashed-dotted line represents the trajectory associated to the average rates $\mathbf{E}_{k \sim \Pi^*}[k]$.

When the density does not accumulate near the boundary, the interplay between Ω , and dynamics of the system and the measured quantities of interest will determine how the optimal density distributes mass in Ω . The experiments in Section 6.2 show that even though this may lead to a better estimate of certain parameters, e.g., the initial condition, the optimal density may assign probability to non-physical values. As mentioned before, this can be mitigated by choosing in (13) a suitable set C of constraints for the expected value. It is interesting to see that some choices of Ω may be better suited for some tasks. As an example, our results in Section 6.2 show that to estimate the initial condition with the expected value it is better to allow for wider intervals. In contrast, to capture the relation between the capacities it may be more appropriate to allow for narrower intervals.

In our experiments we have emphasized the role of the expected value of the initial conditions to assess the performance of the method. However, our results provide a closed-form expression for the density from which other statistics could be evaluated. As an example, in the experiments in Section 6.1 a much better estimate would be the mode. In fact, Fig. 4 suggests that when the system and the measurements are linear, the mode will be at the boundary of Ω . The experiments in Section 6.2 suggest the same. In general, the choice of a suitable statistic would depend on the properties of the system and the available measurements of quantities of interest.

Finally, our results show that the structure of the family of densities that can be found by the method depends strongly on how complex the dynamics of the system are, or how complex are the observations about the system. For linear ODEs and linear observations our results show that the probability density ρ^* must be the positive part of an affine function. As a consequence, to add expressiveness to the probability measure Π^* found by the method it may be necessary to use more complex reference measures μ_Ω on Ω or mixtures.

8. Conclusion

In this work, we propose a Bayesian model to quantify input uncertainty on random differential equations. Furthermore, we construct a family of priors for the probability density of the initial condition for which the *maximum a posteriori* estimate for the density can be computed using Tikhonov regularization on moment constraints for the density. This provides an underlying statistical model for the method proposed by Meyers et al., built on early work by Banks et al., to perform uncertainty quantification in aggregate data problems. Under a convexity assumption on the model, we leveraged duality to deduce an equivalent finite-dimensional formulation, which for typical cases becomes an unconstrained convex problem with a smooth objective function. This allowed us to provide a full theoretical characterization of the family of densities that can be obtained by the method, along with their dependence in the parameters of the problem. Although the trade-off is that the objective function involves a high-dimensional integral, we have shown that the problem can be solved efficiently using standard approximation methods. We believe that our results may lead to the development of more efficient implementations by using tailored numerical integration methods. Furthermore, our analysis yields insight about the flexibility of the approach when performing uncertainty quantification for high-dimensional systems by showing the interplay between the support of the density, the dynamics of the system, and the constraints being enforced.

Our work leaves some questions open. First, our method shows that there is a subtle interplay between the dynamics of the system, the observations about it, and the support of the probability measure that is to be found. Understanding this interplay would help practitioners make an informed choice of support or constraints for the expected values depending on the properties of the system. Second, we have shown that the structure of the probability densities found by the method is restricted by the system and the measured quantities of interest. A question of interest is to determine the size of the class of densities with this structure. Finally, it is an open question whether the optimal probability density can be *refined* to improve the performance of the method according to some criterion. These questions will be explored as future lines of research.

CRediT authorship contribution statement

Elena Villalón: Writing – original draft, Software, Methodology, Conceptualization. **Qian Yang:** Software, Resources, Methodology, Data curation. **Carlos A. Sing Long:** Writing – review & editing, Writing – original draft, Supervision, Software, Methodology, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

We thank the reviewers for their comments and suggestions. C. A. Sing Long thanks A. Jara and M. Castro for their insightful comments and discussions on Bayesian nonparametric methods. C. A. Sing Long and E. Villalón were partially supported by the grant ANID – FONDECYT – 1211643. C. A. Sing Long was also partially supported by the grants ANID – Millennium Science Initiative

Program – NCN17_059 and ANID – Millennium Science Initiative Program – NCN17_129, and by the National Center for Artificial Intelligence CENIA FB210017, Basal ANID. Q. Yang acknowledges support from the National Science Foundation (NSF) Grant No. DMR-2102406.

Appendix A. Proof of main results in Section 2

A.1. Proof of Proposition 1

Define for $i \in \{1, \dots, n\}$ the averages

$$\bar{\varphi}_i = \int_{\Omega} \varphi_i(x) d\mu_{\Omega}(x)$$

and the vector $\bar{\varphi} \in \mathbb{R}^n$ accordingly. Let

$$C_n := \bigcap_{x \in \Omega} \{a \in \mathbb{R}^n : \langle a, \bar{\varphi} \rangle = 1, \langle a, \varphi(x) \rangle \geq 0\}.$$

It is clear this is a closed and convex set in \mathbb{R}^n . If $\rho \in P_{\mu, n}$ then $\rho = \langle a, \varphi \rangle$ for some $a \in \mathbb{R}^n$. Since $\rho \geq 0$ and its integral with respect to μ_{Ω} equals 1 we deduce that

$$\langle a, \bar{\varphi} \rangle = 1 \quad \text{and} \quad x \in \Omega : \langle a, \varphi(x) \rangle \geq 0$$

whence $a \in C_n$. Conversely, if $a \in C_n$ then $\rho := \langle a, \varphi \rangle$ is such that $\rho \geq 0$ then it is straightforward to see that $\rho \in P_{\mu, n}$.

From the hypotheses it follows that $\bar{\varphi}_1 > 0$ whence $\bar{\varphi}_1^{-1} \varphi_1 \in P_{\mu, n}$. Thus, both $P_{\mu, n}$ and C_n are non-empty sets. Let $a_n := \bar{\varphi} / \|\bar{\varphi}\|_2^2$. Then $\langle a_n, \bar{\varphi} \rangle = 1$. Let $v_1, \dots, v_{n-1} \in \mathbb{R}^n$ be orthonormal vectors such that $\langle v_i, \bar{\varphi} \rangle = 0$ for $i \in \{1, \dots, n-1\}$, and let V be the matrix with v_1, \dots, v_{n-1} as its columns. By construction we have that $V^* a_n = 0$. We define the set

$$\Theta_{V_n} := \left\{ \theta \in \mathbb{R}^{n-1} : a_n + \sum_{i=1}^{n-1} \theta_i v_i \in C_n \right\}$$

which is closed and convex since it is the preimage of the affine map $\theta \mapsto a_n + V\theta$. If $\theta \in \Theta_{V_n}$ then, by construction, $a_n + V\theta \in C_n$. Conversely, if $a \in C_n$ we can always write

$$a = a_0 a_n + \sum_{i=1}^{n-1} \theta_i v_i,$$

for some $\theta \in \mathbb{R}^{n-1}$. It suffices to show that $a_0 = 1$. Note that

$$1 = \langle \bar{\varphi}, a \rangle = a_0 \langle \bar{\varphi}, a_n \rangle = a_0$$

whence $\theta \in \Theta_{V_n}$. It remains to show that Θ_{V_n} has non-empty interior. Define $a_1 = \bar{\varphi}_1^{-1} e_1$. Then $\rho = \langle a_1, \varphi \rangle = \bar{\varphi}_1^{-1} \varphi_1$ belongs to $P_{\mu, n}$ whence $a_1 \in C_n$. Therefore, let

$$a = a_n + V(V^* a_1 + \theta) = a_1 + V\theta$$

for some θ to determine. Define

$$\rho = \langle a, \varphi \rangle = \langle a_1, \varphi \rangle + \langle V\theta, \varphi \rangle = \frac{1}{\bar{\varphi}_1} \varphi_1 + \langle V\theta, \varphi \rangle.$$

Note that ρ has integral equal to one for any choice of θ . Therefore, we must show it is non-negative to conclude that $\rho \in P_{\mu, n}$. Since $\varphi_1, \dots, \varphi_n$ are bounded, there is $B > 0$ such that $\|\varphi(x)\|_2 \leq B$ for $x \in \Omega$. Hence, it suffices to have

$$|\langle V\theta, \varphi(x) \rangle| \leq \|V\theta\|_2 \|\varphi(x)\|_2 \leq B \|\theta\|_2 \leq \bar{\varphi}_1^{-1} \inf_{x \in \Omega} \varphi_1(x) \leq \bar{\varphi}_1^{-1} \varphi_1(x)$$

to conclude that $\rho \geq 0$. Since this is satisfied for any θ such that

$$\|\theta\|_2 \leq \frac{1}{B \bar{\varphi}_1} \inf_{x \in \Omega} \varphi_1(x)$$

and the right-hand side is strictly positive, we conclude that Θ_{V_n} has non-empty interior.

A.2. Proof of Proposition 2

Let $\{\rho^{\star, n}\}_{n \in \mathbb{N}}$ be the sequence of solutions to the finite-dimensional problems (10), and let F denote the objective function, which is independent of n . As L is non-negative, F is non-positive. First, since $V_1 \subset V_n$ for any $n \in \mathbb{N}$ we have that $P_{\mu, 1} \subset P_{\mu, n}$ whence $F(\rho^{\star, 1}) \leq F(\rho^{\star, n})$. Since L is non-negative, this implies that $2F(\rho^{\star, 1}) \geq \lambda \|\rho^{\star, n}\|_{L_{\mu}^2}^2$. Hence, using the fact that $P_{\mu, n} \subset P_{\mu}$, we conclude that the sequence $\{\rho^{\star, n}\}_{n \in \mathbb{N}}$ is a norm-bounded sequence in P_{μ} . By the Banach-Alaoglu theorem, by possibly passing to a

subsequence, we can assume without loss that it converges weakly to a density $\rho^{*,\infty}$. Since P_μ is convex and closed, it is also weakly closed. Thus $\rho^{*,\infty} \in P_\mu$.

Now, let ρ^* be the solution to the infinite-dimensional problem (12) and let ρ_n^* be the projection of ρ^* onto the closed convex set $P_{\mu,n}$. Since the closure of the union of the subspaces $\{V_n\}_{n \in \mathbb{N}}$ is all of $L_\mu^2(\Omega)$, we have that $\{\rho_n^*\}_{n \in \mathbb{N}}$ converges in norm to ρ^* . By construction, we have that $F(\rho_n^*) \leq F(\rho^{*,n}) \leq F(\rho^*)$.

Now, remark that the map

$$\rho \mapsto \int_{\Omega} \rho(\mathbf{x}_0) U(\mathbf{x}_0) d\mu_{\Omega}(\mathbf{x}_0)$$

is strongly (or norm) continuous *and* weakly continuous. Since the L_μ^2 -norm also strongly continuous, we conclude that F is strongly continuous whence

$$F(\rho^*) = \limsup_{n \rightarrow \infty} F(\rho_n^*).$$

On the other hand, the L_μ^2 -norm is also strongly lower semicontinuous and thus weakly sequentially lower semicontinuous [27, Theorem 9.1]. This implies that F is weakly sequentially upper semicontinuous whence

$$\limsup_{n \rightarrow \infty} F(\rho^{*,n}) \leq F(\rho^{*,\infty}).$$

Combining these two facts, we conclude that

$$F(\rho^*) = \limsup_{n \rightarrow \infty} F(\rho_n^*) \leq \limsup_{n \rightarrow \infty} F(\rho^{*,n}) \leq F(\rho^{*,\infty}) \leq F(\rho^*)$$

whence $\rho^{*,\infty}$ is a solution to the infinite-dimensional problem, proving the claim.

Appendix B. Proofs of statements in Section 3

B.1. Proof of Proposition 3

Define the linear map $\mathcal{U} : L_\mu^2(\Omega) \rightarrow \mathbb{R}^q$ as

$$\mathcal{U}(\rho) := \int_{\Omega} U(\mathbf{x}_0) \rho(\mathbf{x}_0) d\mu_{\Omega}(\mathbf{x}_0)$$

which, by (6) and the Cauchy-Schwarz inequality, is bounded. As \mathbb{R}^q is finite-dimensional, \mathcal{U} has closed range. Consider the auxiliary problem

$$\begin{aligned} \text{minimize}_{\rho \in L_\mu^2(\Omega)} \quad & \frac{1}{2} \|\rho\|_{L_\mu^2}^2 + \frac{1}{\lambda} L(\mathbf{A}\mathcal{U}(\rho), \mathbf{y}) \\ \text{subject to} \quad & \rho \geq 0, \quad \langle 1, \rho \rangle_{L_\mu^2} = 1, \quad \mathcal{U}(\rho) \in C. \end{aligned}$$

where 1 denotes the function that is identically equal to one over Ω ; since μ_{Ω} is finite, $1 \in L_\mu^2$. The objective function in the above is strongly convex [27, Definition 10.7] and the feasible set is closed and convex. Hence, there is a unique minimizer ρ^* for the above problem [27, Proposition 11.8]. Let $\bar{\mathbf{u}}^* := \mathcal{U}(\rho^*)$. We claim $(\rho^*, \bar{\mathbf{u}}^*)$ is an optimal solution to (13). If $(\rho, \bar{\mathbf{u}})$ is feasible for (13) then ρ is feasible for the above auxiliary problem and thus

$$\begin{aligned} \frac{1}{2} \|\rho^*\|_{L_\mu^2}^2 + \frac{1}{\lambda} L(\mathbf{A}\bar{\mathbf{u}}^*, \mathbf{y}) &= \frac{1}{2} \|\rho^*\|_{L_\mu^2}^2 + \frac{1}{\lambda} L(\mathbf{A}\mathcal{U}(\rho^*), \mathbf{y}) \\ &\leq \frac{1}{2} \|\rho\|_{L_\mu^2}^2 + \frac{1}{\lambda} L(\mathbf{A}\mathcal{U}(\rho), \mathbf{y}) \\ &= \frac{1}{2} \|\rho\|_{L_\mu^2}^2 + \frac{1}{\lambda} L(\mathbf{A}\bar{\mathbf{u}}, \mathbf{y}) \end{aligned}$$

where we used the fact that $\bar{\mathbf{u}} = \mathcal{U}(\rho)$ by feasibility of $(\rho, \bar{\mathbf{u}})$. Hence, the objective value of $(\rho, \bar{\mathbf{u}})$ is at least that of $(\rho^*, \bar{\mathbf{u}}^*)$. We conclude $(\rho^*, \bar{\mathbf{u}}^*)$ is an optimal solution to (13). To show it is the only solution, if we let $(\rho, \bar{\mathbf{u}})$ be any other solution, then, from the strong convexity of the quadratic term on ρ in the objective, it follows that $\rho = \rho^*$. By feasibility, we conclude that $\bar{\mathbf{u}} = \mathcal{U}(\rho) = \mathcal{U}(\rho^*) = \bar{\mathbf{u}}^*$.

B.2. Proof of Theorem 1

First, it is apparent that $\mathcal{B} : X_P \rightarrow X_D$ is bounded and that its adjoint $\mathcal{B}^* : X_D \rightarrow X_P$ is given by

$$\mathcal{B}^*(\eta, \xi, \omega) = \begin{bmatrix} -(\eta + \langle \omega, U(\mathbf{x}_0) \rangle_2) \\ \omega - \mathbf{A}^* \xi \\ \xi \end{bmatrix}$$

The function $f : X_P \rightarrow (-\infty, +\infty]$ is f is proper, convex and lower semicontinuous, with domain

$$\mathbf{dom}(f) = \{(\rho, \bar{\mathbf{u}}, \mathbf{v}) \in X_P : \rho \geq 0, \bar{\mathbf{u}} \in C\},$$

whereas $g : X_D \rightarrow (-\infty, +\infty]$ is proper, convex and lower semicontinuous. Furthermore, it is polyhedral. Note that

$$\mathbf{dom}(g) = \{\mathbf{r}\}$$

and that $\mathbf{dom}(g) \cap \text{relint}(\mathcal{B}(\mathbf{dom}(f))) \neq \emptyset$. In fact,

$$\mathcal{B}(\mathbf{dom}(f)) = \{(s, \mathbf{w}, \mathbf{z}) \in \mathbb{R} \times \mathbb{R}^q \times \mathbb{R}^m : s < 0\}.$$

This follows from the fact that we may choose

$$\rho \equiv \frac{|s|}{\mu_\Omega(\Omega)}, \quad \bar{\mathbf{u}} = \frac{|s|}{\mu_\Omega(\Omega)} \int \mathbf{U}(\mathbf{x}_0) d\mu_\Omega(\mathbf{x}_0) + \mathbf{w} \quad \text{and} \quad \mathbf{v} = \frac{|s|}{\mu_\Omega(\Omega)} \int \mathbf{A}\mathbf{U}(\mathbf{x}_0) d\mu_\Omega(\mathbf{x}_0) + \mathbf{A}\mathbf{w} + \mathbf{z}$$

to obtain $(s, \mathbf{w}, \mathbf{z}) = \mathcal{B}(\rho, \bar{\mathbf{u}}, \mathbf{v})$. By [27, Fact 15.25] we conclude that the origin belongs to the strong relative interior of $\mathbf{dom}(g) - \mathcal{B}(\mathbf{dom}(f))$. Therefore, by [27, Theorem 15.23] strong duality holds. On one hand, the convex dual of f is

$$\begin{aligned} f^*(\rho, \bar{\mathbf{u}}, \mathbf{v}) &= \sup_{(\rho', \bar{\mathbf{u}}', \mathbf{v}') \in X_P} \langle (\rho, \bar{\mathbf{u}}, \mathbf{v}), (\rho', \bar{\mathbf{u}}', \mathbf{v}') \rangle - f(\rho', \bar{\mathbf{u}}', \mathbf{v}') \\ &= \sup_{\rho' \in L_\mu^2} \left(\langle \rho, \rho' \rangle_{L_\mu^2} - \frac{1}{2} \|\rho'\|_{L_\mu^2}^2 - \mathbb{I}_{L_\mu^{2+}}(\rho) \right) \\ &\quad + \sup_{\bar{\mathbf{u}}' \in \mathbb{R}^q} \left(\langle \bar{\mathbf{u}}, \bar{\mathbf{u}}' \rangle_2 - \mathbb{I}_C(\bar{\mathbf{u}}') \right) + \sup_{\mathbf{v}' \in \mathbb{R}^m} \left(\langle \mathbf{v}, \mathbf{v}' \rangle_2 - \frac{1}{\lambda} L(\mathbf{v}', \mathbf{y}) \right) \\ &= \frac{1}{2} \|\rho\|_{L_\mu^2}^2 + \sup_{\rho' \in L_\mu^2} \left(-\frac{1}{2} \|\rho - \rho'\|_{L_\mu^2}^2 - \mathbb{I}_{L_\mu^{2+}}(\rho) \right) + \mathbb{I}_C^*(\bar{\mathbf{u}}) + \frac{1}{\lambda} L^*(\lambda \mathbf{v}, \mathbf{y}) \\ &= \frac{1}{2} \|\rho\|_{L_\mu^2}^2 - \frac{1}{2} \|\rho_-\|_{L_\mu^2}^2 + \mathbb{I}_C^*(\bar{\mathbf{u}}) + \frac{1}{\lambda} L^*(\lambda \mathbf{v}, \mathbf{y}) \\ &= \frac{1}{2} \|\rho_+\|_{L_\mu^2}^2 + \mathbb{I}_C^*(\bar{\mathbf{u}}) + \frac{1}{\lambda} L^*(\lambda \mathbf{v}, \mathbf{y}) \end{aligned}$$

where ρ_+, ρ_- denote the positive and negative parts of ρ . On the other, the convex dual of g is

$$g^*(\eta, \omega, \xi) = \langle (\eta, \omega, \xi), \mathbf{r} \rangle = -\eta.$$

We conclude that the dual objective becomes

$$\frac{1}{2} \int_{\Omega} (\eta + \langle \omega, \mathbf{U}(\mathbf{x}_0) \rangle_2)_+^2 d\mu_\Omega(\mathbf{x}_0) + \mathbb{I}_C^*(\mathbf{A}^* \xi - \omega) + \frac{1}{\lambda} L^*(-\lambda \xi, \mathbf{y}) - \eta$$

as desired. Finally, by [27, Theorem 19.1], if $(\rho^*, \bar{\mathbf{u}}^*, \mathbf{v}^*) \in X_P$ and $(\eta^*, \omega^*, \xi^*) \in X_D$ are optimal primal and dual variables then

$$-\mathcal{B}^*(\eta^*, \omega^*, \xi^*) \in \partial f(\rho^*, \bar{\mathbf{u}}^*, \mathbf{v}^*) \quad \text{and} \quad (\eta^*, \omega^*, \xi^*) \in \partial g(\mathcal{B}(\rho^*, \bar{\mathbf{u}}^*, \mathbf{v}^*)).$$

From

$$\partial f(\rho, \bar{\mathbf{u}}, \mathbf{v}) = \{(\rho + \varphi, 0, \mathbf{g}) : \varphi \leq 0, \rho\varphi = 0, \mathbf{g} \in \lambda^{-1} \partial L(\mathbf{v}, \mathbf{y})\}$$

it follows that the first condition implies

$$\eta^* + \langle \omega^*, \mathbf{U}(\mathbf{x}_0) \rangle = \rho^*(\mathbf{x}_0) + \varphi^*(\mathbf{x}_0).$$

This is possible only if

$$\rho^*(\mathbf{x}_0) = (\eta^* + \langle \omega^*, \mathbf{U}(\mathbf{x}_0) \rangle)_+.$$

B.3. Proof of Theorem 2

This result is proven by slightly modifying the arguments used in the proof of Theorem 1. We use the same notation as in that proof. Define the space $X_P^n := V_n \times \mathbb{R}^q \times \mathbb{R}^m$. Consider the restriction $\mathcal{B} : X_P^n \rightarrow X_D$. The critical difference is to observe that the adjoint $\mathcal{B}^* : X_D \rightarrow X_P^n$ becomes

$$\mathcal{B}^*(\eta, \xi, \omega) = \begin{bmatrix} -(\eta 1_{V_n} + \langle \omega, U_{V_n}(x_0) \rangle_2) \\ \omega - A^* \xi \\ \xi \end{bmatrix}.$$

Define the convex function $f_n : X_P^n \rightarrow \mathbb{R}$

$$f_n(\rho, \bar{u}, v) = \frac{1}{2} \|\rho\|_{L_\mu^2}^2 + \mathbb{I}_{V_n^+}(\rho) + \frac{1}{\lambda} L(v, y) + \mathbb{I}_C(\bar{u})$$

Since there exists a strictly positive function $\varphi_1 \in V_n$ we have that V_n^+ is non-empty. Therefore, the primal problem associated to the finite-dimensional problem is

$$\underset{(\rho, \bar{u}, v) \in X_P^n}{\text{minimize}} \quad f_n(\rho, \bar{u}, v) + g(\mathcal{B}(\rho, \bar{u}, v) - r) \quad (\text{B.1})$$

with dual

$$\underset{(\eta, \omega, \xi) \in X_D^n}{\text{maximize}} \quad -f_n^*(-\mathcal{B}^*(\eta, \omega, \xi)) - g^*(\eta, \omega, \xi). \quad (\text{B.2})$$

To prove strong duality holds, note that f_n is proper, convex and lower semicontinuous, with domain

$$\text{dom}(f_n) = \{(\rho, \bar{u}, v) \in X_P^n : \rho \geq 0, \bar{u} \in C\}.$$

As in the proof of Theorem 1, we have that $\text{dom}(g) \cap \text{relint}(\mathcal{B}(\text{dom}(f_n))) \neq \emptyset$. In fact,

$$\mathcal{B}(\text{dom}(f_n)) = \{(s, w, z) \in \mathbb{R} \times \mathbb{R}^q \times \mathbb{R}^m : s < 0\}$$

follows from the fact that we may choose

$$\rho \equiv \frac{|s|}{\varphi_1} \varphi_1, \quad \bar{u} = \frac{|s|}{\mu_\Omega(\Omega)} \int_{\Omega} U(x_0) d\mu_\Omega(x_0) + w \quad \text{and} \quad v = \frac{|s|}{\mu_\Omega(\Omega)} \int_{\Omega} A U(x_0) d\mu_\Omega(x_0) + Aw + z$$

to obtain $(s, w, z) = \mathcal{B}(\rho, \bar{u}, v)$. Hence, strong duality holds. It remains to compute f_n^* . We have that

$$\begin{aligned} f_n^*(\rho, \bar{u}, v) &= \sup_{(\rho', \bar{u}', v') \in X_P^n} \langle (\rho, \bar{u}, v), (\rho', \bar{u}', v') \rangle - f(\rho', \bar{u}', v') \\ &= \sup_{\rho' \in V_n} \left(\langle \rho, \rho' \rangle_{L_\mu^2} - \frac{1}{2} \|\rho'\|_{L_\mu^2}^2 - \mathbb{I}_{V_n^+}(\rho) \right) \\ &\quad + \sup_{\bar{u}' \in \mathbb{R}^q} \left(\langle \bar{u}, \bar{u}' \rangle_2 - \mathbb{I}_C(\bar{u}') \right) + \sup_{v' \in \mathbb{R}^m} \left(\langle v, v' \rangle_2 - \frac{1}{\lambda} L(v', y) \right) \\ &= \frac{1}{2} \|\rho\|_{L_\mu^2}^2 + \sup_{\rho' \in V_n} \left(-\frac{1}{2} \|\rho - \rho'\|_{L_\mu^2}^2 - \mathbb{I}_{V_n^+}(\rho) \right) + \mathbb{I}_C^*(\bar{u}) + \frac{1}{\lambda} L^*(\lambda v, y) \\ &= \frac{1}{2} \|\rho\|_{L_\mu^2}^2 - \frac{1}{2} \|\rho - \rho'\|_{L_\mu^2}^2 + \mathbb{I}_C^*(\bar{u}) + \frac{1}{\lambda} L^*(\lambda v, y) \\ &= \frac{1}{2} \|\rho_+\|_{L_\mu^2}^2 + \mathbb{I}_C^*(\bar{u}) + \frac{1}{\lambda} L^*(\lambda v, y), \end{aligned}$$

proving the theorem.

B.4. Proof of Proposition 4

Let $h(s) = s_+^2/2$. Then h is differentiable with Lipschitz derivative $h'(s) = s_+$; its Lipschitz constant is 1. Fix x_0 and let

$$q(\alpha, \beta) = \alpha + \langle \beta, U(x_0) \rangle_2.$$

Then

$$\begin{aligned} &\frac{|h(q(\alpha, \beta)) - h(q(\alpha_0, \beta_0)) - h'(q(\alpha_0, \beta_0))(q(\alpha, \beta) - q(\alpha_0, \beta_0))|}{((\alpha - \alpha_0)^2 + \|\beta - \beta_0\|_2^2)^{1/2}} \\ &= \frac{|q(\alpha, \beta) - q(\alpha_0, \beta_0)|}{((\alpha - \alpha_0)^2 + \|\beta - \beta_0\|_2^2)^{1/2}} \left| \int_0^1 h'((1-s)q(\alpha_0, \beta_0) + sq(\alpha, \beta)) - h'(q(\alpha_0, \beta_0)) ds \right| \\ &\leq \frac{1}{2} \frac{(q(\alpha, \beta) - q(\alpha_0, \beta_0))^2}{((\alpha - \alpha_0)^2 + \|\beta - \beta_0\|_2^2)^{1/2}} \end{aligned}$$

$$\begin{aligned} &\leq \frac{(\alpha - \alpha_0)^2 + \|\mathbf{U}(\mathbf{x}_0)\|_2^2 \|\boldsymbol{\beta} - \boldsymbol{\beta}_0\|_2^2}{((\alpha - \alpha_0)^2 + \|\boldsymbol{\beta} - \boldsymbol{\beta}_0\|_2^2)^{1/2}} \\ &\leq \max\{1, \|\mathbf{U}(\mathbf{x}_0)\|_2^2\} ((\alpha - \alpha_0)^2 + \|\boldsymbol{\beta} - \boldsymbol{\beta}_0\|_2^2)^{1/2}. \end{aligned}$$

Hence,

$$\begin{aligned} &\frac{|H(\alpha, \boldsymbol{\beta}) - H(\alpha_0, \boldsymbol{\beta}_0) - \partial_\alpha H(\alpha_0, \boldsymbol{\beta}_0)(\alpha - \alpha_0) - \langle \nabla_{\boldsymbol{\beta}} H(\alpha_0, \boldsymbol{\beta}_0), \boldsymbol{\beta} - \boldsymbol{\beta}_0 \rangle_2|}{((\alpha - \alpha_0)^2 + \|\boldsymbol{\beta} - \boldsymbol{\beta}_0\|_2^2)^{1/2}} \\ &\leq \left(\int_{\Omega} \max\{1, \|\mathbf{U}(\mathbf{x}_0)\|_2^2\} d\mu_{\Omega}(\mathbf{x}_0) \right) ((\alpha - \alpha_0)^2 + \|\boldsymbol{\beta} - \boldsymbol{\beta}_0\|_2^2)^{1/2} \end{aligned}$$

from where the first statement follows. The second statement follows from similar arguments, and the fact that h' is Lipschitz continuous. We omit the details for brevity.

B.5. Proof of Theorem 3

The theorem is a consequence of a change of variables using T . Although it is invertible with continuous inverse, it is not differentiable and it is not Lipschitz over \bar{B}_2 . For this reason, we cannot apply the change of variables formula directly.

We proceed as follows. Let $\delta > 0$. Define

$$A^\delta := \bar{B}_2 \setminus \delta B_2 = \{\mathbf{x} \in \bar{B}_2 : \delta \leq \|\mathbf{x}\|_2 \leq 1\}$$

and

$$\Omega_0^\delta := \{T(\mathbf{x}) : \mathbf{x} \in A^\delta\} = \{\mathbf{x} \in \Omega_0 : \delta \leq \gamma_0(\mathbf{x}) \leq 1\}.$$

Both $\{A^\delta\}_{\delta>0}$ and $\{\Omega_0^\delta\}_{\delta>0}$ are decreasing sequences. Consider the approximation

$$H^\delta(\alpha, \boldsymbol{\beta}) = \frac{1}{2} \int_{\Omega_0^\delta} (\alpha + \langle \boldsymbol{\beta}, \mathbf{U}(\mathbf{x}_0 + \mathbf{x}_c) \rangle_2)_+^2 d\mathbf{x}_0.$$

The integrand is non-negative and $H^\delta \rightarrow H$ pointwise as $\delta \rightarrow 0$ by the monotone convergence theorem [51, Theorem 2.4.1]. We will use the change of variables formula for each H^δ and then conclude by applying a limit argument.

For simplicity, we write $q(\mathbf{x}) = \|\mathbf{x}\|_2 / \gamma_0(\mathbf{x})$. Note q is well-defined for $\mathbf{x} \neq 0$ and there exists $C_q > 0$ such that

$$\forall \mathbf{x} \neq 0 : \frac{1}{C_q} \leq q(\mathbf{x}) \leq C_q.$$

We first prove the following auxiliary result.

Lemma 1. *The map $T : A^\delta \rightarrow \Omega_0^\delta$ defined in (20) is Lipschitz with Lipschitz inverse $T^{-1} : \Omega_0^\delta \rightarrow A^\delta$ defined in (21).*

Proof of Lemma 1. Let $\mathbf{x}, \mathbf{x}' \in A^\delta$. Then

$$T(\mathbf{x}) - T(\mathbf{x}') = q(\mathbf{x})(\mathbf{x} - \mathbf{x}') + (q(\mathbf{x}) - q(\mathbf{x}'))\mathbf{x}'.$$

The first term can be bounded by $\|\mathbf{x} - \mathbf{x}'\|_2$. For the second term, note that

$$\begin{aligned} |q(\mathbf{x}) - q(\mathbf{x}')| &\leq \frac{|\|\mathbf{x}\|_2 - \|\mathbf{x}'\|_2|}{\gamma_0(\mathbf{x})} + \|\mathbf{x}'\|_2 \left| \frac{1}{\gamma_0(\mathbf{x})} - \frac{1}{\gamma_0(\mathbf{x}')} \right| \\ &\leq \frac{1}{\delta} \|\mathbf{x} - \mathbf{x}'\|_2 + \left| \frac{1}{\gamma_0(\mathbf{x})} - \frac{1}{\gamma_0(\mathbf{x}')} \right|. \end{aligned}$$

Finally, note that for $t, s \in [\delta, 1]$ with $s \leq t$

$$\left| \frac{1}{t} - \frac{1}{s} \right| = \left| \int_s^t \frac{d\alpha}{\alpha^2} \right| \leq \frac{1}{\delta^2} |t - s|$$

from where

$$\|T(\mathbf{x}) - T(\mathbf{x}')\|_2 \leq C_q \|\mathbf{x} - \mathbf{x}'\|_2 + \frac{1}{\delta} \|\mathbf{x} - \mathbf{x}'\|_2 + \frac{1}{\delta^2} |\gamma_0(\mathbf{x}) - \gamma_0(\mathbf{x}')|$$

and we conclude from the fact that γ_0 is Lipschitz continuous. To prove this claim, remark that since γ_0 is continuous, it is uniformly continuous on $\bar{B}_2(0, 1)$. Let $\varepsilon > 0$ and let $\delta > 0$ be such that

$$\|\mathbf{x}' - \mathbf{x}\|_2 < \delta \Rightarrow |\gamma_0(\mathbf{x}') - \gamma_0(\mathbf{x})| < \varepsilon$$

for $\mathbf{x}', \mathbf{x} \in \bar{B}_2(0, 1)$. Then, for any $\mathbf{x}', \mathbf{x} \in \mathbb{R}^d$ with $\mathbf{x}' \neq \mathbf{x}$ we can choose $\alpha = \delta/2\|\mathbf{x}' - \mathbf{x}\|_2$ so that

$$|\gamma_0(\mathbf{x}') - \gamma_0(\mathbf{x})| = \frac{1}{\lambda} |\gamma_0(\lambda\mathbf{x}') - \gamma_0(\lambda\mathbf{x})| < \frac{\varepsilon}{\lambda} = \frac{\varepsilon}{\delta} \|\mathbf{x}' - \mathbf{x}\|_2.$$

The same arguments can be applied to prove T^{-1} is Lipschitz continuous. We omit the details for brevity. \square

Hence, T is differentiable Lebesgue a.e. by Rademacher's theorem [52, Theorem 2, Section 3.1.2]. Let DT denote its differential matrix. Then, we can apply the change of variables formula [52, Lemma 1, Section 3.3.1 and Theorem 2, Section 3.3.3], to obtain

$$H^\delta(\alpha, \beta) = \frac{1}{2} \int_{A^\delta} (\alpha + \langle \beta, U(T(\mathbf{x}_0) + \mathbf{x}_c) \rangle_2)_+^2 |\det(DT(\mathbf{x}_0))| d\mathbf{x}_0.$$

We can further simplify the term involving the Jacobian using the following lemma.

Lemma 2. *The function $q : A^\delta \rightarrow \mathbb{R}^p$ is differentiable Lebesgue a.e. Furthermore, if \mathbf{x} is a point of differentiability, then $\langle \nabla q(\mathbf{x}), \mathbf{x} \rangle_2 = 0$.*

Proof of Lemma 2. The arguments in the proof of Lemma 1 show q is Lipschitz continuous. Hence, by Rademacher's theorem it is differentiable Lebesgue a.e. If \mathbf{x} is a point of differentiability, then

$$\lim_{\mathbf{x}' \rightarrow \mathbf{x}} \frac{|q(\mathbf{x}') - q(\mathbf{x}) - \langle \nabla q(\mathbf{x}), \mathbf{x}' - \mathbf{x} \rangle_2|}{\|\mathbf{x}' - \mathbf{x}\|_2} = 0.$$

However, for any $\alpha > 0$ and $\mathbf{x} \neq 0$

$$q(\alpha\mathbf{x}) = \frac{\|\alpha\mathbf{x}\|_2}{\gamma_0(\alpha\mathbf{x})} = \frac{|\alpha|\|\mathbf{x}\|_2}{|\alpha|\gamma_0(\mathbf{x})} = q(\mathbf{x})$$

whence q is homogeneous of degree 0. If we take $\mathbf{x}' = (1 + \alpha)\mathbf{x}$ for α sufficiently small, it follows that $\mathbf{x}' - \mathbf{x} = \alpha\mathbf{x}$ and

$$\lim_{\alpha \rightarrow 0} \frac{|\alpha| |\langle \nabla q(\mathbf{x}), \mathbf{x} \rangle_2|}{|\alpha| \|\mathbf{x}\|_2} = \frac{|\langle \nabla q(\mathbf{x}), \mathbf{x} \rangle_2|}{\|\mathbf{x}\|_2} = 0$$

whence the lemma follows. \square

Therefore, there is a subset of A^δ of full Lebesgue measure where both T and q are differentiable. If \mathbf{x} is such a point, then

$$DT(\mathbf{x}) = q(\mathbf{x})\mathbf{I}_d + \mathbf{x}\nabla q(\mathbf{x})^t \Rightarrow \det(DT(\mathbf{x})) = q(\mathbf{x})^d (1 + q(\mathbf{x})^{-d} \langle \nabla q(\mathbf{x}), \mathbf{x} \rangle_2) = q(\mathbf{x})^d$$

where we used the identity $\det(\mathbf{I} + \bar{\mathbf{u}}\mathbf{v}^t) = 1 + \langle \bar{\mathbf{u}}, \mathbf{v} \rangle_2$. We conclude that

$$H^\delta(\alpha, \beta) = \frac{1}{2} \int_{A^\delta} (\alpha + \langle \beta, U(q(\mathbf{x})\mathbf{x} + \mathbf{x}_c) \rangle_2)_+^2 q(\mathbf{x})^d d\mathbf{x}.$$

As $H^\delta \rightarrow H$ pointwise, the theorem follows.

Appendix C. Proof of main results in Section 4

C.1. Proof of Proposition 5

We use the same notation as in the proof of Proposition 5. Note that

$$R_U(\bar{\mathbf{u}}) = \inf \left\{ \frac{1}{2} \|\rho\|_{L_\mu^2}^2 : \rho \in L_\mu^2(\Omega), \rho \geq 0, \langle 1, \rho \rangle_{L_\mu^2} = 1, \mathcal{U}(\rho) = \bar{\mathbf{u}} \right\}. \quad (\text{C.1})$$

Hence, $R_U(\bar{\mathbf{u}}) = +\infty$ if and only if the set

$$\{\rho \in L_\mu^2(\Omega) : \rho \geq 0, \langle 1, \rho \rangle_{L_\mu^2} = 1, \mathcal{U}(\rho) = \bar{\mathbf{u}}\}$$

is empty. Hence, the domain of R_U is given by (24). It is direct to verify it is convex. To show it is non-empty, it suffices to choose $\rho \equiv 1/\mu_\Omega(\Omega)$ to conclude that

$$\frac{1}{\mu_\Omega(\Omega)} \int_{\Omega} U(\mathbf{x}_0) d\mu_\Omega(\mathbf{x}_0) \in \text{dom}(R_U)$$

whence R_U is proper.

To prove the convexity of R_U , observe that the objective in (C.1) is strongly convex [27, Definition 10.7] and that the feasible set is closed. Then, there exists a unique ρ_u such that [27, Proposition 11.8]

$$R_U(\bar{\mathbf{u}}) = \frac{1}{2} \|\rho_u\|_{L_\mu^2}^2.$$

Let $\bar{\mathbf{u}}_1, \bar{\mathbf{u}}_2 \in \mathbf{dom}(R_U)$, let $\rho_1 = \rho_{u_1}$ and $\rho_2 = \rho_{u_2}$, and let $\theta \in [0, 1]$. Since

$$\theta\rho_1 + (1-\theta)\rho_2 \geq 0, \quad \langle 1, \theta\rho_1 + (1-\theta)\rho_2 \rangle_{L_\mu^2} = 1 \quad \text{and} \quad \mathcal{U}(\theta\rho_1 + (1-\theta)\rho_2) = \theta\bar{\mathbf{u}}_1 + (1-\theta)\bar{\mathbf{u}}_2$$

we see that

$$\begin{aligned} R_U(\theta\bar{\mathbf{u}}_1 + (1-\theta)\bar{\mathbf{u}}_2) &\leq \frac{1}{2} \|\theta\rho_1 + (1-\theta)\rho_2\|_{L_\mu^2}^2 \\ &\leq \frac{1}{2} \theta \|\rho_1\|_{L_\mu^2}^2 + \frac{1}{2} (1-\theta) \|\rho_2\|_{L_\mu^2}^2 \\ &= \theta R_U(\bar{\mathbf{u}}_1) + (1-\theta) R_U(\bar{\mathbf{u}}_2). \end{aligned}$$

Therefore, we conclude R_U is convex.

C.2. Proof of Proposition 6

By the choice of \mathbf{U} we have $\mathbf{dom}(R_U) \subset \Omega$. To show $\mathbf{int}(\Omega) \subset \mathbf{dom}(R_U)$ let $\mu_0 \in \mathbf{int}(\Omega)$ and let $r > 0$ be such that $B(\mu_0, r) \subset \mathbf{int}(\Omega)$. Then, it suffices to choose

$$\rho(\mathbf{x}_0) := \frac{1}{\mu(B(\mu_0, r))} \chi_{B(\mu_0, r)}$$

to conclude that $\mu_0 \in \mathbf{dom}(R_U)$. Now, it remains to show that $\mu_0 \in \mathbf{bd}(\Omega) \notin \mathbf{dom}(R_U)$. Suppose there is $\rho \in L_\mu^2(\Omega)$ with $\rho \geq 0$ and $\langle 1, \rho \rangle_{L_\mu^2} = 1$ such that

$$\mu_0 = \int_{\Omega} \rho(\mathbf{x}_0) \mathbf{x}_0 d\mathbf{x}_0 \Rightarrow \int_{\Omega} \rho(\mathbf{x}_0) (\mathbf{x}_0 - \mu_0) d\mathbf{x}_0 = 0.$$

Let $\mathbf{z} \in N_{\Omega}(\mu_0)$ where N_{Ω} denotes the normal cone. Define for $s \geq 0$

$$\Omega_s := \{\mathbf{x}_0 \in \Omega : \langle \mathbf{z}, \mathbf{x}_0 - \mu_0 \rangle \leq s\}$$

and observe that $\Omega = \Omega_0$. If $\rho \geq 0$ on a measurable set $A \subset \Omega_s$ for some $s < 0$ then

$$\int_A \rho(\mathbf{x}_0) \langle \mathbf{z}, \mathbf{x}_0 - \mu_0 \rangle_2 d\mathbf{x}_0 \leq s \int_A \rho(\mathbf{x}_0) d\mathbf{x}_0 < 0.$$

However, this implies

$$\begin{aligned} 0 &= \int_{\Omega \cap A} \rho(\mathbf{x}_0) \langle \mathbf{z}, \mathbf{x}_0 - \mu_0 \rangle_2 d\mathbf{x}_0 + \int_{\Omega \setminus A} \rho(\mathbf{x}_0) \langle \mathbf{z}, \mathbf{x}_0 - \mu_0 \rangle_2 d\mathbf{x}_0 \\ &< s \int_A \rho(\mathbf{x}_0) d\mathbf{x}_0 + \int_{\Omega \setminus A} \rho(\mathbf{x}_0) \langle \mathbf{z}, \mathbf{x}_0 - \mu_0 \rangle_2 d\mathbf{x}_0. \end{aligned}$$

We conclude that

$$0 < -s \int_A \rho(\mathbf{x}_0) d\mathbf{x}_0 < \int_{\Omega \setminus A} \rho(\mathbf{x}_0) \langle \mathbf{z}, \mathbf{x}_0 - \mu_0 \rangle_2 d\mathbf{x}_0 \leq 0$$

which is a contradiction unless $s = 0$. Therefore, ρ must be supported on

$$\Omega \cap \{\mathbf{x}_0 \in \mathbb{R}^d : \langle \mathbf{z}, \mathbf{x}_0 - \mu_0 \rangle_2 = 0\} \subset \{\mathbf{x}_0 \in \mathbb{R}^d : \langle \mathbf{z}, \mathbf{x}_0 - \mu_0 \rangle_2 = 0\}.$$

Since the set in the right-hand side has Lebesgue measure zero, we conclude $\rho = 0$. This contradiction proves the first statement of the proposition. To prove the second, let $\mu_0 \in \mathbf{bd}(\Omega)$ and let $\{\mu_n\}_{n \in \mathbb{N}}$ be a sequence in $\mathbf{int}(\Omega)$ with $\mu_n \rightarrow \mu_0$. Let $\{\rho_n\}_{n \in \mathbb{N}}$ be the sequence of optimal solutions to (14) for each μ_n . Suppose that $\{R_U(\mu_n)\}_{n \in \mathbb{N}}$ remains bounded by a constant $C > 0$. This implies that $\|\rho_n\|_{n \in \mathbb{N}}$ is norm bounded in $L^2(\Omega)$. Since norm bounded sequences are weakly convergent the sequence has a weak limit $\rho_0 \in L^2(\Omega)$ [27, Theorem 2.33]. Since the set $\{\rho \in L^2(\Omega) : \rho \geq 0, \langle 1, \rho \rangle_{L^2} = 1\}$ is norm closed, it is weakly closed and ρ_0 also belongs to this set [27, Theorem 3.34]. However, the function $f(\mathbf{x}_0) = x_{0,i}$ is in $L^2(\Omega)$ for each $i \in \{1, \dots, d\}$. Therefore,

$$\int_{\Omega} \rho_n(\mathbf{x}_0) x_{0,i} d\mathbf{x}_0 = \mu_n \rightarrow \mu_0 = \int_{\Omega} \rho_0(\mathbf{x}_0) x_{0,i} d\mathbf{x}_0.$$

However, μ_0 lies in the boundary. Using the same arguments as before, we reach a contradiction. This proves the claim.

C.3. Proof of Proposition 7

We show that R_U° is invariant under transformations of the form $T(\boldsymbol{\mu}) = \mathbf{Q}(\boldsymbol{\mu} - \boldsymbol{\mu}_0) + \boldsymbol{\mu}_0$ for any $d \times d$ orthogonal matrix \mathbf{Q} . The value of $R_U^\circ(T(\mathbf{x}_0))$ is the optimal value of the problem

$$\begin{aligned} & \underset{\rho \in L_\mu^2(\Omega): \rho \geq 0}{\text{minimize}} \quad \frac{1}{2} \int_{\Omega} \rho(\mathbf{x}_0)^2 d\mathbf{x}_0 \\ & \text{subject to} \quad \int_{\Omega} \rho(\mathbf{x}_0) d\mu_{\Omega}(\mathbf{x}_0) = 1, \quad \int_{\Omega} \rho(\mathbf{x}_0) \mathbf{x}_0 d\mu_{\Omega}(\mathbf{x}_0) = T(\boldsymbol{\mu}). \end{aligned}$$

However, T induces an isometric isomorphism $\mathcal{T} : L_\mu^2(\Omega) \mapsto L_\mu^2(\Omega)$ given by $\mathcal{T}(\rho)(\mathbf{x}_0) = \rho(T(\mathbf{x}_0))$. Therefore, the optimal value to the above problem is equivalent to

$$\begin{aligned} & \underset{\rho \in L_\mu^2(\Omega): \rho \geq 0}{\text{minimize}} \quad \frac{1}{2} \int_{\Omega} \rho(\mathbf{x}_0)^2 d\mathbf{x}_0 \\ & \text{subject to} \quad \int_{\Omega} \rho(\mathbf{x}_0) d\mathbf{x}_0 = 1, \quad \int_{\Omega} \rho(\mathbf{x}_0) T(\mathbf{x}_0) d\mathbf{x}_0 = T(\boldsymbol{\mu}), \end{aligned}$$

whence $R_U^\circ(T(\boldsymbol{\mu})) = R_U^\circ(\boldsymbol{\mu})$. It follows that the function $f : \bar{B}_2(0, 1) \rightarrow \mathbb{R}$ given by $f(\mathbf{z}) = R_U^\circ(r\mathbf{z} + \boldsymbol{\mu}_0)$ is invariant under any orthogonal transformation. We conclude that $f(\mathbf{z}) = \varphi(\|\mathbf{z}\|_2)$ for some non-negative function φ and, by a suitable change of variable, we arrive to the first statement. Note that $\varphi \geq 0$ as $R_U^\circ \geq 0$. Furthermore, since $f(\mathbf{z}) = \varphi(\|\mathbf{z}\|_2)$ is convex on $\bar{B}_2(0, 1)$ then φ must be convex and non-decreasing. Finally, from Proposition 6 we deduce that $\varphi(s) \rightarrow \infty$ as $s \rightarrow 1$ as this implies $\boldsymbol{\mu} \rightarrow \text{bd}(\Omega)$. The second statement follows by standard arguments and the fact that φ is non-decreasing. We omit the proof for brevity.

Appendix D. Proof of main results in Section 5

D.1. Proof of Proposition 8

We have the bound

$$\begin{aligned} \sigma^2(\boldsymbol{\alpha}, \boldsymbol{\beta}) &= \int_{\Omega} (h(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}_0) - H(\boldsymbol{\alpha}, \boldsymbol{\beta}))^2 d\Pi_{\Omega}(\mathbf{x}_0) \\ &= \int_{\Omega} \left(\int_{\Omega} (h(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}_0) - h(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}'_0)) d\Pi_{\Omega}(\mathbf{x}'_0) \right)^2 d\Pi_{\Omega}(\mathbf{x}_0) \\ &\leq \iint_{\Omega} (h(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}_0) - h(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}'_0))^2 d\Pi_{\Omega}(\mathbf{x}'_0) d\Pi_{\Omega}(\mathbf{x}_0). \end{aligned}$$

From the identity

$$\frac{1}{2}(c + x_2)_+^2 - \frac{1}{2}(c + x_1)_+^2 = (x_2 - x_1) \int_0^1 (c + \theta x_2 + (1 - \theta)x_1)_+ d\theta$$

for any $c, x_1, x_2 \in \mathbb{R}$ we deduce

$$\begin{aligned} |h(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}_0) - h(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}'_0)| &\leq \left(\int_0^1 (\alpha + \langle \boldsymbol{\beta}, (1 - \theta)\mathbf{U}(\mathbf{x}'_0) + \theta\mathbf{U}(\mathbf{x}_0) \rangle)_+ d\theta \right) |\langle \boldsymbol{\beta}, \mathbf{U}(\mathbf{x}_0) - \mathbf{U}(\mathbf{x}'_0) \rangle_2| \\ &\leq \left(\int_0^1 ((1 - \theta)(\alpha + \langle \boldsymbol{\beta}, \mathbf{U}(\mathbf{x}'_0) \rangle_2)_+ + \theta(\alpha + \langle \boldsymbol{\beta}, \mathbf{U}(\mathbf{x}_0) \rangle_2)_+) d\theta \right) |\langle \boldsymbol{\beta}, \mathbf{U}(\mathbf{x}_0) - \mathbf{U}(\mathbf{x}'_0) \rangle_2| \\ &= \left(\frac{1}{2}(\alpha + \langle \boldsymbol{\beta}, \mathbf{U}(\mathbf{x}'_0) \rangle_2)_+ + \frac{1}{2}(\alpha + \langle \boldsymbol{\beta}, \mathbf{U}(\mathbf{x}_0) \rangle_2)_+ \right) |\langle \boldsymbol{\beta}, \mathbf{U}(\mathbf{x}_0) - \mathbf{U}(\mathbf{x}'_0) \rangle_2| \end{aligned}$$

$$\leq \|\beta\|_2 \left(|\alpha| + \frac{1}{2} \|\beta\|_2 (\|U(x_0)\|_2 + \|U(x'_0)\|_2) \right) \|U(x_0) - U(x'_0)\|_2$$

the proposition follows.

D.2. Proof of Proposition 9

Since

$$g_n(\alpha, \beta) - \nabla H(\alpha, \beta) = \frac{1}{M} \sum_{m=1}^M \iint_{\Omega} (\nabla h(\alpha, \beta, x_0^{(n,m)}) - \nabla h(\alpha, \beta, x_0)) d\Pi_{\Omega}(x_0) d\Pi_{\Omega}(x_0^{(n,m)})$$

we have

$$\begin{aligned} \mathbf{E}_{x_0 \sim \Pi_{\Omega}} \|g_n(\alpha, \beta) - \nabla H(\alpha, \beta)\|_2^2 &= \frac{1}{M^2} \sum_{m=1}^M \iint_{\Omega} \|\nabla h(\alpha, \beta, x_0^{(n,m)}) - \nabla h(\alpha, \beta, x_0)\|_2^2 d\Pi_{\Omega}(x_0) d\Pi_{\Omega}(x_0^{(n,m)}) \\ &= \frac{1}{M} \iint_{\Omega} \|\nabla h(\alpha, \beta, x'_0) - \nabla h(\alpha, \beta, x_0)\|_2^2 d\Pi_{\Omega}(x_0) d\Pi_{\Omega}(x'_0) \end{aligned}$$

From

$$\begin{aligned} |\partial_{\alpha} h(\alpha, \beta, x'_0) - \partial_{\alpha} h(\alpha, \beta, x_0)| &= |(\alpha + \langle \beta, U(x'_0) \rangle_2)_+ - (\alpha + \langle \beta, U(x_0) \rangle_2)_+| \\ &\leq \|\beta\|_2 \|U(x'_0) - U(x_0)\|_2 \end{aligned}$$

and

$$\begin{aligned} \|\nabla_{\beta} h(\alpha, \beta, x'_0) - \nabla_{\beta} h(\alpha, \beta, x_0)\|_2 &\leq |(\alpha + \langle \beta, U(x'_0) \rangle_2)_+ - (\alpha + \langle \beta, U(x_0) \rangle_2)_+| \|U(x'_0)\|_2 \\ &\quad + (\alpha + \langle \beta, U(x_0) \rangle_2)_+ \|U(x'_0) - U(x_0)\|_2 \\ &\leq (\|\beta\|_2 \|U(x'_0)\|_2 + |\alpha| + \|\beta\|_2) \|U(x'_0) - U(x_0)\|_2 \end{aligned}$$

we deduce

$$\|\nabla h(\alpha, \beta, x'_0) - \nabla h(\alpha, \beta, x_0)\|_2^2 \leq 2 \|U(x'_0) - U(x_0)\|_2^2 (\|\beta\|_2^2 + (\|\beta\|_2 \|U(x'_0)\|_2 + |\alpha| + \|\beta\|_2)^2)$$

from where the inequality follows. We omit the details for brevity.

References

- [1] M. Braun, Differential Equations and Their Applications: An Introduction to Applied Mathematics, Texts in Applied Mathematics, vol. 11, Springer, New York, NY, 1993, <http://link.springer.com/10.1007/978-1-4612-4360-1>.
- [2] S. Wiggins, Introduction to Applied Nonlinear Dynamical Systems and Chaos, Texts in Applied Mathematics, vol. 2, Springer-Verlag, New York, 2003, <http://link.springer.com/10.1007/b97481>.
- [3] E. Hairer, G. Wanner, C. Lubich, Geometric Numerical Integration, Springer Series in Computational Mathematics, vol. 31, Springer-Verlag, Berlin/Heidelberg, 2006, <http://link.springer.com/10.1007/3-540-30666-8>.
- [4] A. Iserles, A First Course in the Numerical Analysis of Differential Equations, 2nd edition, Cambridge University Press, 2008, <https://www.cambridge.org/core/product/identifier/9780511995569/type/book>.
- [5] M. Arnold, Numerical methods for simulation in applied dynamics, in: G. Maier, J. Salençon, W. Schneider, B. Schrefler, P. Serafini, M. Arnold, W. Schiehlen (Eds.), Simulation Techniques for Applied Dynamics, in: CISM International Centre for Mechanical Sciences., vol. 507, Springer, Vienna, Vienna, 2008, pp. 191–246, http://link.springer.com/10.1007/978-3-211-89548-1_5.
- [6] H.-G. Beyer, B. Sendhoff, Robust optimization – a comprehensive survey, Comput. Methods Appl. Mech. Eng. 196 (33–34) (2007) 3190–3218, <https://doi.org/10.1016/j.cma.2007.03.003>, <https://linkinghub.elsevier.com/retrieve/pii/S0045782507001259>.
- [7] C.J. Roy, W.L. Oberkampf, A comprehensive framework for verification, validation, and uncertainty quantification in scientific computing, Comput. Methods Appl. Mech. Eng. 200 (25–28) (2011) 2131–2144, <https://doi.org/10.1016/j.cma.2011.03.016>, <https://linkinghub.elsevier.com/retrieve/pii/S0045782511001290>.
- [8] R.C. Smith, Uncertainty Quantification: Theory, Implementation, and Applications, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2013, <https://pubs.siam.org/doi/book/10.1137/1.9781611973228>.
- [9] D. Kumar, F. Ahmed, S. Usman, A. Alajo, S.B. Alam, Recent advances in uncertainty quantification methods for engineering problems, in: AI Assurance, Elsevier, 2023, pp. 453–472, <https://linkinghub.elsevier.com/retrieve/pii/B978032391917000275>.
- [10] H.T. Banks, K.B. Flores, I.G. Rosen, E.M. Rutter, M. Sirlancı, W.C. Thompson, The Prohorov metric framework and aggregate data inverse problems for random PDEs, Commun. Appl. Anal. 22 (3) (2018) 415–446.
- [11] C. Schacht, A. Meade, H. Banks, H. Enderling, D. Abate-Daga, Estimation of probability distributions of parameters using aggregate population data: analysis of a CAR T-cell cancer model, Math. Biosci. Eng. 16 (6) (2019) 7299–7326, <https://doi.org/10.3934/mbe.2019365>, <http://www.aimspress.com/article/10.3934/mbe.2019365>.
- [12] H. Banks, A.E. Meade, C. Schacht, J. Catenacci, W.C. Thompson, D. Abate-Daga, H. Enderling, Parameter estimation using aggregate data, Appl. Math. Lett. 100 (2020) 105999, <https://doi.org/10.1016/j.aml.2019.105999>, <https://linkinghub.elsevier.com/retrieve/pii/S0893965919303234>.
- [13] H.T. Banks, S. Hu, W.C. Thompson, H.T. Banks, Modeling and Inverse Problems in the Presence of Uncertainty, Monographs and Research Notes in Mathematics, CRC Press/Chapman & Hall, Boca Raton, Fla., 2014.

- [14] H.T. Banks, W. Botsford, F. Kappel, C. Wang, Modeling and estimation in size-structured population models, in: 2nd Course on Mathematical Ecology, Trieste, December 8–12, 1986, World Press, Singapore, 1988, pp. 521–541.
- [15] H.T. Banks, H.T. Banks, A Functional Analysis Framework for Modeling, Estimation and Control in Science and Engineering, A Chapman & Hall book CRC Press, Chapman & Hall, Boca Raton, Fla., 2012.
- [16] J. Meyers, J. Rogers, A. Gerlach, Koopman operator method for solution of generalized aggregate data inverse problems, *J. Comput. Phys.* 428 (2021) 110082, <https://doi.org/10.1016/j.jcp.2020.110082>, <https://linkinghub.elsevier.com/retrieve/pii/S0021999120308561>.
- [17] A.R. Gerlach, A. Leonard, J. Rogers, C. Rackauckas, The Koopman expectation: an operator theoretic method for efficient analysis and optimization of uncertain hybrid dynamical systems, arXiv:2008.08737 [math], <http://arxiv.org/abs/2008.08737>, Aug. 2020.
- [18] A.N. Tikhonov, V.Y. Arsenin, *Solutions of Ill-Posed Problems*, Halsted Press, New York, 1977.
- [19] X. Jiang, S. Li, R. Furfaro, Z. Wang, Y. Ji, High-dimensional uncertainty quantification for Mars atmospheric entry using adaptive generalized polynomial chaos, *Aerosp. Sci. Technol.* 107 (2020) 106240, <https://doi.org/10.1016/j.ast.2020.106240>, <https://linkinghub.elsevier.com/retrieve/pii/S1270963820309226>.
- [20] I. Abdallah, C. Lataniotis, B. Sudret, Parametric hierarchical Kriging for multi-fidelity aero-servo-elastic simulators — application to extreme loads on wind turbines, *Probab. Eng. Mech.* 55 (2019) 67–77, <https://doi.org/10.1016/j.probengmech.2018.10.001>, <https://linkinghub.elsevier.com/retrieve/pii/S0266892017302126>.
- [21] Z. He, Z. Zhang, High-dimensional uncertainty quantification via tensor regression with rank determination and adaptive sampling, *IEEE Trans. Compon. Packag. Manuf. Technol.* 11 (9) (2021) 1317–1328, <https://doi.org/10.1109/TCMT.2021.3093432>, <https://ieeexplore.ieee.org/document/9467329>.
- [22] P. Hall, On Kullback-Leibler loss and density estimation, *Ann. Stat.* 15 (4) (Dec. 1987), <https://doi.org/10.1214/aos/1176350606>, <https://projecteuclid.org/journals/annals-of-statistics/volume-15/issue-4/On-Kullback-Leibler-Loss-and-Density-Estimation/10.1214/aos/1176350606.full>.
- [23] S.N. MacEachern, Nonparametric Bayesian methods: a gentle introduction and overview, *Commun. Stat. Appl. Methods* 23 (6) (2016) 445–466, <https://doi.org/10.5351/CSAM.2016.23.6.445>, <http://www.csam.or.kr/journal/view.html?doi=10.5351/CSAM.2016.23.6.445>.
- [24] B.O. Koopman, Hamiltonian systems and transformation in Hilbert space, *Proc. Natl. Acad. Sci.* 17 (5) (1931) 315–318, <https://doi.org/10.1073/pnas.17.5.315>, <https://pnas.org/doi/full/10.1073/pnas.17.5.315>.
- [25] M. Budišić, R. Mohr, I. Mezić, Applied koopmanism, chaos: an interdisciplinary, *J. Nonlinear Sci.* 22 (4) (2012) 047510, <https://doi.org/10.1063/1.4772195>, <https://pubs.aip.org/cha/article/22/4/047510/341880/Applied-Koopmanism>.
- [26] R.T. Rockafellar, R.J.B. Wets, *Variational Analysis*, Grundlehren der Mathematischen Wissenschaften, vol. 317, Springer, Berlin Heidelberg, 1998, <http://link.springer.com/10.1007/978-3-642-02431-3>.
- [27] H.H. Bauschke, P.L. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, CMS Books in Mathematics, Springer, New York, NY, 2011, <https://link.springer.com/10.1007/978-1-4419-9467-7>.
- [28] Y. Nesterov, A method for solving the convex programming problem with convergence rate $\mathcal{O}(1/k^2)$, *Dokl. Akad. Nauk SSSR* 269 (3) (1983) 543–547.
- [29] Y. Nesterov, Gradient methods for minimizing composite functions, *Math. Program.* 140 (1) (2013) 125–161, <https://doi.org/10.1007/s10107-012-0629-5>, <http://link.springer.com/10.1007/s10107-012-0629-5>.
- [30] A. Beck, M. Teboulle, A fast iterative shrinkage-thresholding algorithm for linear inverse problems, *SIAM J. Imaging Sci.* 2 (1) (2009) 183–202, <https://doi.org/10.1137/080716542>, <http://pubs.siam.org/doi/10.1137/080716542>.
- [31] A.M. Rubinov, A.A. Yagubov, The space of star-shaped sets and its applications in nonsmooth optimization, in: M.L. Balinski, R.W. Cottle, L.C.W. Dixon, B. Korte, M.J. Toddla, E.L. Allgower, W.H. Cunningham, J.E. Dennis, B.C. Eaves, R. Fletcher, D. Goldfarb, J.-B. Hiriart-Urruty, M. Iri, R.G. Jeroslow, D.S. Johnson, C. Lemarechal, L. Lovasz, L. McLinden, M.J.D. Powell, W.R. Pulleyblank, A.H.G. Rinnooy Kan, K. Ritter, R.W.H. Sargent, D.F. Shanno, L.E. Trotter, H. Tuy, R.J.B. Wets, E.M.L. Beale, G.B. Dantzig, L.V. Kantorovich, T.C. Koopmans, A.W. Tucker, P. Wolfe, V.F. Demyanov, L.C.W. Dixon (Eds.), *Quasidifferential Calculus*, in: Mathematical Programming Studies, vol. 29, Springer, Berlin Heidelberg, 1986, pp. 176–202, <http://link.springer.com/10.1007/BFb0121146>.
- [32] A.S. Nemirovskii, D.B. Yudin, *Problem Complexity and Method Efficiency in Optimization*, Wiley-Interscience Series in Discrete Mathematics, Wiley Chichester, New York, 1983.
- [33] A.H. Stroud, D. Secrest, *Gaussian Quadrature Formulas*, Prentice-Hall, Englewood Cliffs, NJ, 1966.
- [34] J. Glaubitz, Constructing positive interpolatory cubature formulas, arXiv:2009.11981 (Sep. 2020), <http://arxiv.org/abs/2009.11981>.
- [35] T. Gerstner, M. Griebel, Numerical integration with sparse grids, *Numer. Algorithms* 18 (3/4) (1998) 209–232, <https://doi.org/10.1023/A:1019129717644>, <http://link.springer.com/10.1023/A:1019129717644>.
- [36] G. Petrova, Cubature formulae for spheres, simplices and balls, *J. Comput. Appl. Math.* 162 (2) (2004) 483–496, <https://doi.org/10.1016/j.cam.2003.08.036>, <https://linkinghub.elsevier.com/retrieve/pii/S0377042703007611>.
- [37] D. Williams, L. Shunn, A. Jameson, Symmetric quadrature rules for simplexes based on sphere close packed lattice arrangements, *J. Comput. Appl. Math.* 266 (2014) 18–38, <https://doi.org/10.1016/j.cam.2014.01.007>, <https://linkinghub.elsevier.com/retrieve/pii/S0377042714000211>.
- [38] S.A. Smolyak, Quadrature and interpolation formulas for tensor products of certain classes of functions, *Dokl. Akad. Nauk SSSR* 148 (5) (1963) 1042–1045.
- [39] J. Lu, D.L. Darmofal, Higher-dimensional integration with Gaussian weight for applications in probabilistic design, *SIAM J. Sci. Comput.* 26 (2) (2004) 613–624, <https://doi.org/10.1137/S1064827503426863>, <http://pubs.siam.org/doi/10.1137/S1064827503426863>.
- [40] I. Mysovskikh, The approximation of multiple integrals by using interpolatory cubature formulae, in: *Quantitative Approximation*, Elsevier, 1980, pp. 217–243, <https://linkinghub.elsevier.com/retrieve/pii/B9780122136504500258>.
- [41] J.H. Halton, On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals, *Numer. Math.* 2 (1) (1960) 84–90, <https://doi.org/10.1007/BF01386213>, <http://link.springer.com/10.1007/BF01386213>.
- [42] H. Niederreiter, Low-discrepancy and low-dispersion sequences, *J. Number Theory* 30 (1) (1988) 51–70, [https://doi.org/10.1016/0022-314X\(88\)90025-X](https://doi.org/10.1016/0022-314X(88)90025-X), <https://linkinghub.elsevier.com/retrieve/pii/0022314X8890025X>.
- [43] N.A. Butler, A new class of equal-weight integration rules on the hypercube, *Numer. Math.* 99 (2) (2004) 349–363, <https://doi.org/10.1007/s00211-004-0562-5>, <http://link.springer.com/10.1007/s00211-004-0562-5>.
- [44] A. Nemirovski, A. Juditsky, G. Lan, A. Shapiro, Robust stochastic approximation approach to stochastic programming, *SIAM J. Optim.* 19 (4) (2009) 1574–1609, <https://doi.org/10.1137/070704277>, <http://pubs.siam.org/doi/10.1137/070704277>.
- [45] W.J. Morokoff, R.E. Caflisch, Quasi-Monte Carlo integration, *J. Comput.* 122 (2) (1995) 218–230, <https://doi.org/10.1006/jcph.1995.1209>, <https://linkinghub.elsevier.com/retrieve/pii/S0021999185712090>.
- [46] J. Brauchart, E. Saff, I. Sloan, R. Womersley, QMC designs: optimal order quasi Monte Carlo integration schemes on the sphere, *Math. Comput.* 83 (290) (2014) 2821–2851, <https://doi.org/10.1090/S0025-5718-2014-02839-1>, <https://www.ams.org/mcom/2014-83-290/S0025-5718-2014-02839-1>.
- [47] R.L. Smith, Efficient Monte Carlo procedures for generating points uniformly distributed over bounded regions, *Oper. Res.* 32 (6) (1984) 1296–1308.
- [48] L. Lovász, S. Vempala, Hit-and-run from a corner, *SIAM J. Comput.* 35 (4) (2006) 985–1005, <https://doi.org/10.1137/S009753970544727X>, <http://pubs.siam.org/doi/10.1137/S009753970544727X>.
- [49] D.J. Higham, Modeling and simulating chemical reactions, *SIAM Rev.* 50 (2) (2008) 347–368, <https://doi.org/10.1137/060666457>, <http://pubs.siam.org/doi/10.1137/060666457>.
- [50] Q. Yang, C.A. Sing-Long, E.J. Reed, Rapid data-driven model reduction of nonlinear dynamical systems including chemical reaction networks using ℓ^1 -regularization, *Chaos* 30 (5) (2020) 053122, <https://doi.org/10.1063/1.5139463>, <https://pubs.aip.org/aip/cha/article/341976>.
- [51] D.L. Cohn, *Measure Theory*, Birkhäuser Advanced Texts Basler Lehrbücher, second edition, Springer, New York, NY, 2013, <https://link.springer.com/10.1007/978-1-4614-6956-8>.
- [52] L.C. Evans, R.F. Gariepy, *Measure Theory and Fine Properties of Functions*, 1st edition, Chapman and Hall/CRC, 2015, <https://www.taylorfrancis.com/books/9781482242393>.