Confidence intervals for functionals in constrained inverse problems via data-adaptive sampling-based calibration

Michael Stanley Pau Batlle[†] Pratik Patil[‡] mcstanle@alumni.cmu.edu pbatllef@caltech.edu pratikpatil@berkeley.edu Houman Owhadi[†] Mikael Kuusela[§]

owhadi@caltech.edu

Mikael Kuusela[§] mkuusela@andrew.cmu.edu

Abstract

We address functional uncertainty quantification for ill-posed inverse problems where it is possible to evaluate a possibly rank-deficient forward model, the observation noise distribution is known, and there are known parameter constraints. We present four constraint-aware confidence interval constructions extending the theoretical test inversion framework in Batlle et al. (2023) by making the intervals both computationally feasible and less conservative. Our approach first shrinks the potentially unbounded constraint set compact in a data-adaptive way, obtains samples of the relevant test statistic inside this set to estimate a quantile function, and then uses these computed quantities to produce the constraint-aware confidence intervals. Our approach to bounding the constraint set in a data-adaptive way is based on the approach by Berger and Boos (1994), and involves defining a subset of the constraint set where the true parameter is guaranteed to exist with high probability. The probabilistic guarantee of this compact subset is then incorporated into the final coverage guarantee in the form of an uncertainty budget. We then propose custom sampling algorithms to efficiently sample from this subset, even when the parameter space is high-dimensional. Optimization-based interval methods formulate confidence interval computation as two endpoint optimizations, where the optimization constraints can be set to achieve different types of interval calibration while seamlessly incorporating parameter constraints. However, choosing optimization constraints to obtain coverage for a single functional has been elusive. We show that all four proposed intervals achieve nominal coverage for a particular functional both theoretically and in practice, with several numerical examples demonstrating superior performance of our intervals over the OSB interval in terms of both coverage and expected interval length. In particular, we show the superior performance of our intervals in a realistic unfolding simulation from high-energy physics that is severely ill-posed and involves a rank-deficient forward model.

[†]Department of Computing and Mathematical Sciences, California Institute of Technology, CA 91125, USA. [‡]Department of Statistics, University of California, Berkeley, CA 94720, USA.

[§]Department of Statistics and Data Science, Carnegie Mellon University, Pittsburgh, PA 15213, USA.

1 Introduction

This paper proposes a novel uncertainty quantification (UQ) approach for ill-posed inverse problems characterized by a known parametric forward model¹ $f : \mathbb{R}^p \to \mathbb{R}^n$ mapping a parameter $\boldsymbol{x} \in \mathbb{R}^p$ to an observation $f(\boldsymbol{x}) \in \mathbb{R}^n$, additive noise with a known distribution $\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$, and constraints on the forward model parameters denoted by $\boldsymbol{x} \in \mathcal{X}$, where \mathcal{X} is primarily assumed to be of the form $\boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{b}$. We consider scenarios where the parameter space dimension p may exceed the observation space dimension n, which often results in an ill-posed problem. Additionally, we do not require the forward model f to be injective, allowing for the possibility that f is many-to-one. We assume observations are generated by $\boldsymbol{y} = f(\boldsymbol{x}^*) + \boldsymbol{\varepsilon}$, where $\boldsymbol{x}^* \in \mathcal{X}$ is the true but unknown model parameter. Our UQ object of interest is a confidence interval², $\mathcal{C}_{\alpha}(\boldsymbol{y})$, on a one-dimensional quantity of interest (QoI) derived from \boldsymbol{x}^* by $\boldsymbol{\varphi} : \mathcal{X} \to \mathbb{R}$. Our approach produces constraint-aware intervals with a finite-sample frequentist coverage guarantee, namely, for each $\alpha \in (0, 1)$

$$\mathbb{P}(\varphi(\boldsymbol{x}^*) \in \mathcal{C}_{\alpha}(\boldsymbol{y})) \ge 1 - \alpha \text{ for all } \boldsymbol{x}^* \in \mathcal{X}.$$
(1)

The probability in the coverage guarantee (1) is taken over the distribution of the observation $\boldsymbol{y} = f(\boldsymbol{x}^*) + \boldsymbol{\varepsilon}$, where $\boldsymbol{\varepsilon}$ is the additive noise. The level $\alpha \in (0, 1)$ corresponds to the desired coverage probability, meaning that the interval $C_{\alpha}(\boldsymbol{y})$ will contain the true value $\varphi(\boldsymbol{x}^*)$ with probability at least $1 - \alpha$ for any realization of the noise $\boldsymbol{\varepsilon}$.

This paper builds on the confidence interval framework of Batlle et al. (2023), which posits a more general data-generating process, $\boldsymbol{y} \sim P_{\boldsymbol{x}^*}$, and proposes a novel test-inversion framework for constructing constraint-aware confidence intervals for inverse problem QoI's. By proposing a testing framework to evaluate if $\boldsymbol{x}^* \in \Phi_{\mu} := \{\boldsymbol{x} \in \mathcal{X} : \varphi(\boldsymbol{x}) = \mu\}$ via the test statistic $\lambda(\mu, \boldsymbol{y})$, Batlle et al. (2023) construct a confidence set as follows (see Equation (2.9)),

$$\mathcal{C}_{\alpha}(\boldsymbol{y}) := \{ \mu : \lambda(\mu, \boldsymbol{y}) \le q_{\alpha}(\mu) \}.$$
(2)

As shown in Lemma 2.2 of Batlle et al. (2023), there are two ways to set $q_{\alpha}(\mu)$ using a supremum over a particular quantile function, $Q_{\boldsymbol{x}}(1-\alpha)$ (we show a particular realization of such a quantile function in the left-most panel of Figure 1.1), such that $\mathcal{C}_{\alpha}(\boldsymbol{y})$ achieves the desired coverage,

$$Q_{\mu,1-\alpha}^{\max} := \sup_{\boldsymbol{x} \in \Phi_{\mu} \cap \mathcal{X}} Q_{\boldsymbol{x}}(1-\alpha), \tag{3}$$

$$Q_{1-\alpha}^{\max} := \sup_{\boldsymbol{x} \in \varphi(\mathcal{X})} Q_{\mu,1-\alpha}^{\max} = \sup_{\boldsymbol{x} \in \mathcal{X}} Q_{\boldsymbol{x}}(1-\alpha).$$
(4)

In practice, not only are Equations (3) and (4) difficult to compute because $Q_{\boldsymbol{x}}(1-\alpha)$ is typically only accessible via sampling, but the optimizations can be expressed equivalently as chanceconstrained programs (see Lemma 3.1 in Batlle et al. (2023)), which are typically challenging to solve. Furthermore, although $Q_{\mu,1-\alpha}^{\max}$ and $Q_{1-\alpha}^{\max}$ are optimal with respect to the setup in which they are analyzed, they can still produce overly conservative (larger than necessary) confidence sets if the true \boldsymbol{x}^* is far away from the worst-case \boldsymbol{x} against which the above maximum quantiles protect. This paper addresses these practical and statistical challenges by including a bounded data-informed subset of the parameter space, \mathcal{B}_{η} with $\eta \in (0, 1)$, such that $\mathbb{P}(\boldsymbol{x}^* \in \mathcal{B}_{\eta}) \geq 1 - \eta$ (see Section 3.1). We refer to \mathcal{B}_{η} as the "Berger-Boos" set due to its inspiration from Berger and Boos

¹Even though we write the forward model as a known function f, we only require evaluation access to the map throughout the paper.

 $^{^{2}}$ We interchangeably use the terms "confidence set" and "confidence interval" since the latter may be obtained by the former by simply retaining the set endpoints.



Figure 1.1: (Left) A particular quantile surface, $Q_{\boldsymbol{x}}(1-\alpha)$, where $\boldsymbol{x} \geq \boldsymbol{0}$ and $\alpha = 0.32$. This surface was obtained via Monte Carlo sampling the LLR test statistic over a grid of \boldsymbol{x} 's defined by the two-dimensional constrained-Gaussian scenario similar to that in Section 6.1, but with $\boldsymbol{h} := \begin{pmatrix} 1 & 1 \end{pmatrix}^{\top}$. (Right-Top) An illustration of the Berger-Boos set and other, a $1-\eta$ confidence set for \boldsymbol{x}^* , which prevents having to contend with an unbounded parameter space. Additionally, it can eliminate "worst-case" parameter settings in \mathcal{X} that are far from the data, thus potentially making the resulting intervals less conservative. (Right-Bottom) An illustration of an LLR test statistic curve, $\lambda(\mu, \boldsymbol{y})$, over the QoI domain and the interval endpoints that result from using either Equation (5) for the "Sliced" interval or Equation (6) for the "Global" interval.

(1994). We show that we can reformulate Equations (3) and (4) to be constrained over $\mathcal{B}_{\eta} \subset \mathcal{X}$ if we also use a slightly larger quantile level, $1 - \gamma$ for $\gamma \in (\alpha, 1)$, and such that $1 - \alpha = (1 - \eta)(1 - \gamma)$, producing the following maximum quantiles,

$$\bar{q}_{\gamma,\eta}(\mu) := \sup_{\boldsymbol{x} \in \Phi_{\mu} \cap \mathcal{B}_{\eta}} Q_{\boldsymbol{x}}(1-\gamma), \tag{5}$$

$$\bar{q}_{\gamma,\eta} := \sup_{\boldsymbol{x} \in \mathcal{B}_{\eta}} Q_{\boldsymbol{x}}(1-\gamma).$$
(6)

Since $\bar{q}_{\gamma,\eta}(\mu)$ and $\bar{q}_{\gamma,\eta}$ are defined over bounded subsets of \mathcal{X} , we can propose a novel sampling procedure to obtain design points covering \mathcal{B}_{η} and use quantile regression to estimate the desired quantile surface, $Q_{\boldsymbol{x}}(1-\gamma)$. The estimated quantile surface is then used to produce estimates of either $\bar{q}_{\gamma,\eta}(\mu)$ or $\bar{q}_{\gamma,\eta}$ and constructs a confidence set using a set definition similar to Equation (2),

$$C_{\alpha}(\boldsymbol{y}; \mathcal{B}_{\eta}) := \{ \mu : \lambda(\mu, \boldsymbol{y}) \le q \},$$
(7)

where q is set to one of $\bar{q}_{\gamma,\eta}(\mu)$ or $\bar{q}_{\gamma,\eta}$ to obtain the desired interval. We present a variety of ways to perform the sampling and compare the different confidence set formulations and computations.

1.1 Context and related work

With advances in data collection and computational processing, high-dimensional ill-posed inverse problems have become more prevalent, especially in fields like remote sensing and data assimilation. This setting includes a wide array of physical science applications, spanning Earth science (Rodgers, 2000), atmospheric science and remote sensing (Liu et al., 2016; Patil et al., 2022), and high energy physics (Kuusela, 2016; Stanley et al., 2022), among many others. Providing guarantees for UQ in parameter inference from indirect observations is essential for assessing the precision of scientific inferences made in these contexts. However, the inherent ill-posed nature of these problems often leads to inferences that are highly sensitive to noise, posing significant challenges for UQ. Namely, the ill-posedness leads to an identifiability issue in which statistical inference is impossible without providing some form of regularization, which usually takes either a deterministic (e.g., SVD truncation (Höcker and Kartvelishvili, 1996) and Tikhonov regularization (Schmitt, 2012)) or probabilistic (e.g., priors and Bayesian inference) form. Under some assumptions, these two approaches are mathematically equivalent and are therefore subject to the same pitfalls, such as a disruption of the coverage guarantees of downstream intervals as a result of the incurred regularization bias as thoroughly discussed in Kuusela (2016). Our method's focus on one-dimensional QoI's and incorporation of parameter constraints allows for implicit regularization and therefore produces intervals with the promised coverage guarantee while avoiding the problem of regularization bias altogether.

Although including parameter constraints provides implicit regularization and enables handling non-trivial null spaces, it shifts complexity to statistical inference under constraints, a non-trivial problem in even elementary settings (Gouriéroux et al., 1982; Wolak, 1987; Robertson et al., 1988; Shapiro, 1988; Wolak, 1989; Molenberghs and Verbeke, 2007). One elegant solution to this problem originates in what we refer to as *optimization-based UQ*, originating in the work of Burrus et al. (1965) and Rust and Burrus (1972). Stark (1992a) extended and generalized their approach, calling this method *strict bounds*, since it produced guaranteed simultaneous coverage interval estimators complying with known physical constraints on the model parameters. This collection of work defines endpoint optimization problems over the physically constrained parameter space to directly compute confidence intervals for one-dimensional QoI's. Not only does the optimization form of the confidence interval computation shift statistical inference complexity to numerical optimization. but it also allows known physical constraints to be directly included in the endpoint optimizations. This practical advantage is dulled by the difficulty of proving the coverage properties of the intervals resulting from the defined endpoint optimizations in the one-at-a-time setting. Even under the relatively strong assumptions in Rust and Burrus (1972) of a linear forward model, non-negativity parameter constraints and linear QoI, the authors were only able to conjecture the coverage of their interval (known as the Burrus Conjecture). The interval coverage was unsuccessfully proven in Rust and O'Leary (1994) (the error pointed out in Tenorio et al. (2007)) and finally generally refuted in Batlle et al. (2023). These optimized-based intervals that have gained recent attention take the form:

$$\mathcal{I}(\psi_{\alpha}^{2},\boldsymbol{y}) := \left[\varphi^{l}(\psi_{\alpha}^{2},\boldsymbol{y}),\varphi^{u}(\psi_{\alpha}^{2},\boldsymbol{y})\right] = \left[\min_{\boldsymbol{x}\in D(\psi_{\alpha}^{2},\boldsymbol{y})}\varphi(\boldsymbol{x}),\max_{\boldsymbol{x}\in D(\psi_{\alpha}^{2},\boldsymbol{y})}\varphi(\boldsymbol{x})\right],\tag{8}$$

where

$$D(\psi_{\alpha}^{2}, \boldsymbol{y}) := \left\{ \boldsymbol{x} \in \mathcal{X} : \|\boldsymbol{y} - f(\boldsymbol{x})\|_{2}^{2} \le \psi_{\alpha}^{2} \right\}.$$
(9)

The statistical challenge for intervals of this form is the choosing of ψ_{α}^2 such that $\mathcal{I}(\psi_{\alpha}^2, \boldsymbol{y})$ has the desired coverage guarantee. It is important to note that the definition of $D(\psi_{\alpha}^2, \boldsymbol{y})$ assumes that the noise has been standardized; in other words, the covariance matrix of the noise has been transformed to the identity matrix.

The literature proposes two settings to guarantee coverage. One approach is to set ψ_{α}^2 such that $D(\psi_{\alpha}^2, \boldsymbol{y})$ is itself a confidence set for \boldsymbol{x}^* in the parameter space. Since this choice would then

automatically guarantee coverage for $\mathcal{I}(\psi_{\alpha}^2, \boldsymbol{y})$ regardless of the chosen QoI, this setting has been called "simultaneous" in Rust and O'Leary (1994), and "Simultaneous Strict Bounds" (SSB) in Stanley et al. (2022); Batlle et al. (2023) since it also aligns with the setting of the strict bounds construction in Stark (1992a). Under the Gaussian assumption, this can be achieved by setting $\psi_{SSB,\alpha}^2 := \chi_{n,\alpha}^2$, where $\chi_{n,\alpha}^2$ is the upper α -quantile for a chi-squared distribution with *n* degrees of freedom. Although this approach is relatively simple and obtains the desired coverage guarantee, it is conservative since the guarantee holds for all possible QoI choices simultaneously. To tailor the interval to one particular QoI, there is the "one-at-a-time" setting as described in Rust and Burrus (1972); Rust and O'Leary (1994), or "one-at-a-time strict bounds" (OSB) as called in Stanley et al. (2022); Batlle et al. (2023). Under the Gaussian assumption this proposed setting was $\psi_{OSB,\alpha}^2 = \chi_{1,\alpha}^2 + s(\boldsymbol{y})^2$, where $s(\boldsymbol{y})^2 := \min_{\boldsymbol{x} \in \mathcal{X}} \|\boldsymbol{y} - f(\boldsymbol{x})\|_2^2$. Unlike the simultaneous setting, $D(\psi_{OSB,\alpha}^2, \boldsymbol{y})$ is not a $1 - \alpha$ confidence set for \boldsymbol{x}^* , which makes proving its coverage guarantee difficult. The validity of this claim was proposed by the Burrus conjecture (Rust and Burrus, 1972) and generally disproven by Batlle et al. (2023).

To address the challenge of calibrating strict bounds and optimization-based confidence intervals, Batlle et al. (2023) approached these intervals as an inverted likelihood ratio test as shown in Equation (2). The distribution of the log-likelihood ratio (LLR) statistic is non-standard due to the presence of the constraints, adding to the complexity of controlling the type-I error of the test. This framework development allowed for the general disproving of the long-standing Burrus conjecture, invalidating the general calibration of the intervals proposed by Rust and Burrus (1972). However, the framework revealed that the characterization of $Q_x(1-\alpha)$ allowed for proper interval calibration via Equations (3) or (4) in much more general settings than those in which the Burrus conjecture applied (linear forward model, Gaussian noise, and linear functional).

1.2 Our approach and contributions

To address the computational challenges and statistical conservatism following from the use of quantiles (3) or (4) in confidence set (2), this paper develops analogous maximum quantiles over the bounded Berger–Boos set (quantiles (5) and (6)) for the altered confidence set (7). We refer to the confidence intervals produced by quantile (5) as "Sliced", since it considers the maximum quantile along slices of the Berger–Boos set as defined by level-sets of the QoI. We refer to the confidence intervals produced by quantile (6) as "Global", since it considers the maximum quantile over the entire Berger–Boos set. For both the sliced and global forms of the confidence interval, we present and develop both optimization and sampling approaches. In total, these confidence interval construction options produce four different interval varieties.

To validate the coverage guarantees of these interval varieties, we conduct a series of numerical experiments to demonstrate the effectiveness of the intervals in various scenarios. Since the OSB interval described in Patil et al. (2022) and Stanley et al. (2022) is the current standard approach in this setting, it is our primary point of comparison. First, we present low-dimensional examples. The first example is a two-dimensional constrained Gaussian noise model, a well-studied case in the literature (see Tenorio et al. (2007), Batlle et al. (2023)), illustrating how our intervals are competitive with the OSB interval in a scenario where the OSB interval is known to achieve nominal coverage. The second example is a three-dimensional constrained Gaussian noise model, illustrating the advantage of our method in achieving nominal coverage where the OSB interval fails. Second, we apply our method to a simulated version of particle unfolding from high-energy physics (see Stanley et al. (2022)), a binned deconvolution problem, in an 80-dimensional parameter space with a rank-deficient linear forward model. We use the previously studied realistic setting from Stanley

et al. (2022) where the OSB interval is known to provide coverage to show that our Sliced intervals substantially outperform OSB in terms of expected length. Then, we use a parameter setting where the OSB interval does not achieve nominal coverage, but our intervals do, with the Sliced intervals still out-performing OSB in terms of expected length. This application highlights the practical significance of our approach, with our confidence intervals consistently achieving nominal coverage and often outperforming existing methods in terms of expected interval length.

The contributions of this work are both methodological and computational. Methodologically, we propose new confidence interval constructions for the setting described at the beginning of the section. Although the use of the Berger-Boos set is inspired by Berger and Boos (1994), they originally applied it in a hypothesis testing for handling nuisance parameters. Masserano et al. (2024) applied a similar idea to nuisance parameters in a classification setting, but to the best of our knowledge, our paper is the first to apply the idea in the ill-posed inverse problem UQ setting. Inspired by the methods used in simulation-based inference (Dalmasso et al. (2020, 2024); Masserano et al. (2023)) we apply the approach of estimating the quantile function of a test statistic using simulated data. However, this work differs in the underlying model assumptions and composite nature of the null hypotheses. In addition, the prior work assumed bounded and relatively low-dimensional parameter spaces. The search for confidence interval endpoints is sometimes done via the Robbins-Monro (RM) procedure, which iteratively refines estimates to achieve desired coverage probabilities (Garthwaite and Buckland, 1992; Carpenter, 1999). Methods based on quantile regression (Koenker, 2005) have been proposed to improve accuracy. For instance, Fisher et al. (2020) introduced a technique that inverts estimated quantiles to determine the endpoints of the confidence intervals. This work is the first to combine test inversion, sampling and quantile regression to estimate calibrated constraint-aware confidence intervals.

Computationally, we propose two custom sampling methods designed specifically to draw design points within the defined Berger–Boos set. This set can be particularly challenging to sample since it is the intersection between a pre-image based on the observed data and linear parameter constraints. which can produce sharp edges and corners. This set is also elongated due to the ill-posedness of the underlying inverse problem. While there are some approaches to handle similar scenarios arising in more traditional MCMC sampling (e.g., nested sampling (Skilling, 2004; Buchner, 2023) and Hamiltonian Monte Carlo sampling with constraints (Pakman and Paninski, 2014)), these algorithms can be intricate, so we develop the following alternatives. In the low-dimensional setting (p < 10) where the forward model is linear and the noise distribution is Gaussian, we design an accept-reject sampler based on sampling from p-balls as described in Voelker et al. (2017). Although the approach in Voelker et al. (2017) can rapidly sample from the pre-image ellipsoids as defined by the data-generating model assumptions, with even a simple non-negativity constraint on the parameters, the proportion of rejected points quickly get impractically large for even moderate dimensions. For moderate and higher dimensional settings (p > 10), we leverage the Vaidya walk MCMC algorithm from Chen et al. (2018) to generate random walks about a polytope enclosing the Berger–Boos set. Using the Chebyshev ball center to define a notion of polytope centrality along with the extreme points of the Berger–Boos set with respect to the QoI, we create two lines along which starting positions are defined to generate a collection of parallel Markov chains. The design points resulting from these chains are then combined to create the complete set of sampled design points spanning the full range of the QoI values over the Berger–Boos set.

The rest of the paper proceeds as follows. In Section 2, we recapitulate the framework and theoretical components of Batlle et al. (2023) upon which this paper's work is built. Section 3 then presents the theoretical Global and Sliced intervals under the Berger–Boos set formulation, along with

the four interval constructions and theoretical justifications to guarantee that the constructed sampling-based intervals converge in probability to their theoretical counterparts. In Section 4, we give theoretical results regarding our proposed methods. Section 5 presents the custom sampling algorithms along with a brief description of how quantile regression works in our intervals. Section 6 presents four numerical experiments to demonstrate the coverage and length advantages of our intervals over the OSB interval. Finally, Section 7 provides some discussion and conclusions.

2 Background

A key part of the approach in Batlle et al. (2023) was to view the optimization-based intervals shown by Equation (8) as inverted hypothesis tests, more specifically, a particular inverted likelihood ratio test. This change in perspective allowed for interval analysis through the lens of the properties of a particular log-likelihood ratio (LLR) test. We summarize this connection here, as it is critical to our method development carried out in Section 3. Although the numerical examples in Section 6 focus on the linear-Gaussian version of the data-generating process, we follow the more general exposition of Batlle et al. (2023) to indicate that this framework is not limited to that scenario.

Suppose $\boldsymbol{y} \sim P_{\boldsymbol{x}^*}$, where $P_{\boldsymbol{x}^*}$ is a distribution parameterized by a fixed but unknown $\boldsymbol{x}^* \in \mathbb{R}^p$. Let $\ell_{\boldsymbol{x}}(\boldsymbol{y})$ denote the log-likelihood of \boldsymbol{x} evaluated at \boldsymbol{y} . We furthermore suppose we know a set $\mathcal{X} \subseteq \mathbb{R}^p$ such that $\boldsymbol{x}^* \in \mathcal{X}$.

The duality between hypothesis tests and confidence sets is well known in statistics (see, e.g., Chapter 7 of Casella and Berger (2002) or Chapter 5 of Panaretos (2016)). We invert the following hypothesis test:

$$H_0: \boldsymbol{x}^* \in \Phi_\mu \cap \mathcal{X} \quad \text{versus} \quad H_1: \boldsymbol{x}^* \in \mathcal{X} - \Phi_\mu, \tag{10}$$

where $\Phi_{\mu} := \{ \boldsymbol{x} : \varphi(\boldsymbol{x}) = \mu \}$. Since we are looking to obtain a confidence interval on the real line, it makes sense that each hypothesis test is defined by $\mu \in \mathbb{R}$. Notice that the composite structure of the null hypothesis includes all parameter settings within the μ -level set of the functional of interest.

We define the following test statistic to evaluate Test (10):

$$\lambda(\mu, \boldsymbol{y}) := \begin{cases} \inf_{\boldsymbol{x} \in \Phi_{\mu} \cap \mathcal{X}} \left(-2\ell_{\boldsymbol{x}}(\boldsymbol{y})\right) & -\inf_{\boldsymbol{x} \in \mathcal{X}} \left(-2\ell_{\boldsymbol{x}}(\boldsymbol{y})\right), & \text{if } \Phi_{\mu} \cap \mathcal{X} \neq \emptyset, \\ \infty, & \text{otherwise.} \end{cases}$$
(11)

where μ denotes the level set of the null hypothesis, \boldsymbol{y} is the observed data, and \mathcal{X} is the constraint set that is known to contain the parameter. The test is rejected if the test statistic is large, and hence we automatically reject the null if $\Phi_{\mu} \cap \mathcal{X} = \emptyset$ independently of the observed data. Alternatively, we can consider the test (10) only for $\mu \in \varphi(\mathcal{X})$ to perform test inversion. We control the behavior of this test statistic, and therefore the test, by bounding from above the probability of erroneously rejecting the null hypothesis (i.e., type-1 error). As such, we consider the distribution of $\lambda(\mu, \boldsymbol{y})$ under the null. For each $\boldsymbol{x} \in \mathcal{X}$, let $\mu = \varphi(\boldsymbol{x})$. Define $Q_{\boldsymbol{x}} : (0, 1) \to \mathbb{R}$ such that, for all $\alpha \in (0, 1)$,

$$\mathbb{P}\left(\lambda(\mu, \boldsymbol{y}) \le Q_{\boldsymbol{x}}(1-\alpha)\right) = 1 - \alpha,\tag{12}$$

where the probability is over $\boldsymbol{y} \sim P_{\boldsymbol{x}}$. We refer to $Q_{\boldsymbol{x}}$ as the quantile function of the LLR under the null hypothesis at $(\boldsymbol{x}, \varphi(\boldsymbol{x}) = \mu)$. Since the null hypothesis is composite, using this quantile function to define a cutoff is not enough to control type-1 error. Thus, we use the *sliced* maximum quantile function over the level-set under consideration:

$$Q_{\mu,1-\alpha}^{\max} := \sup_{\boldsymbol{x}\in\Phi_{\mu}\cap\mathcal{X}} Q_{\boldsymbol{x}}(1-\alpha).$$
(13)

Note, we refer to this version as "sliced" since we are considering the maximum of the quantile function defined by the level-set of the functional which defines a slice through the constrained parameter space. $Q_{\mu,1-\alpha}^{\max}$ as defined above controls the type-1 error for a specific value of μ . We can define a more conservative confidence set using the following *global* maximum quantile:

$$Q_{1-\alpha}^{\max} := \sup_{\mu \in \varphi(\mathcal{X})} Q_{\mu,1-\alpha}^{\max} = \sup_{\boldsymbol{x} \in \mathcal{X}} Q_{\boldsymbol{x}}(1-\alpha).$$
(14)

By Lemma 2.1 in Batlle et al. (2023), the set

$$\mathcal{C}^{\mu}_{\alpha}(\boldsymbol{y};\mathcal{X}) := \{\mu : \lambda(\mu, \boldsymbol{y}) \le Q^{\max}_{\mu, 1-\alpha}\} \subset \mathbb{R}$$
(15)

defines a $1 - \alpha$ confidence set for the true functional value, $\mu^* = \varphi(\boldsymbol{x}^*)$. The global quantile (14) is more conservative than the sliced quantile (13) since it holds for all null hypotheses, and therefore the set

$$\mathcal{C}_{\alpha}(\boldsymbol{y}; \mathcal{X}) := \{ \mu : \lambda(\mu, \boldsymbol{y}) \le Q_{1-\alpha}^{\max} \} \subset \mathbb{R}$$
(16)

is also a $1 - \alpha$ confidence set.

To connect this framework back to Interval (8), suppose the data generating process is of the form, $P_x = \mathcal{N}(\mathbf{K}x, \mathbf{I})$. The LLR test statistic is then

$$\lambda(\mu, \boldsymbol{y}) = \min_{\boldsymbol{x} \in \Phi_{\mu} \cap \mathcal{X}} \|\boldsymbol{y} - \boldsymbol{K}\boldsymbol{x}\|_{2}^{2} - \min_{\boldsymbol{x} \in \mathcal{X}} \|\boldsymbol{y} - \boldsymbol{K}\boldsymbol{x}\|_{2}^{2}$$
(17)

Batlle et al. (2023) (Theorem 2.4) proved that

$$\left[\inf_{\mu:\lambda(\mu,\boldsymbol{y})\leq Q_{1-\alpha}^{\max}}\mu,\sup_{\mu:\lambda(\mu,\boldsymbol{y})\leq Q_{1-\alpha}^{\max}}\mu\right] = \left[\inf_{\boldsymbol{x}\in D(Q_{1-\alpha}^{\max}+s(\boldsymbol{y})^2,\boldsymbol{y})}\varphi(\boldsymbol{x}),\sup_{\boldsymbol{x}\in D(Q_{1-\alpha}^{\max}+s(\boldsymbol{y})^2,\boldsymbol{y})}\varphi(\boldsymbol{x})\right],$$
(18)

where $D(\cdot, \boldsymbol{y})$ is defined by Equation (9). This equivalence asserts that by setting $\psi_{\alpha}^2 := Q_{1-\alpha}^{\max} + s(\boldsymbol{y})^2$, where $s(\boldsymbol{y})^2 := \min_{\boldsymbol{x} \in \mathcal{X}} \|\boldsymbol{y} - \boldsymbol{K}\boldsymbol{x}\|_2^2$, we guarantee coverage for Interval (8). Furthermore, it asserts that the original OSB interval formulation is only valid if and only if $\chi_{1,\alpha}^2 \ge Q_{1-\alpha}^{\max}$ (see Lemma 2.2 and Corollary 2.3). Batlle et al. (2023) showed that this inequality does not hold in general.

As explored in Batlle et al. (2023) and reiterated above, Interval (8) can be calibrated by computing $Q_{\mu,1-\alpha}^{\max}$ or $Q_{1-\alpha}^{\max}$. However, pursuing calibration in this way is computationally challenging and statistically conservative. Both of these values require the ability to evaluate Q_x . Without parameter constraints, this quantile function can be constant under Gaussian-linear assumptions (Batlle et al., 2023), i.e., the test statistic is pivotal. But even with a relatively simple two-dimensional Gaussian noise model with non-negativity constraints, this quantile function becomes non-trivial (e.g., see Figure 5.3 in Batlle et al. (2023)). Beyond the practical difficulty of dealing with the underlying quantile function, both (13) and (14) can be expressed as chance-constrained optimization problems, which are known to be NP-hard in general and would need to be solved over an unbounded constraint set (Geng and Xie, 2019; Pena-Ordieres et al., 2020; Batlle et al., 2023). Statistically,

since both (13) and (14) are the largest quantile function values subject to their respective constraints, they are by definition conservative, especially in scenarios where the true x^* is far from these most conservative points. Since we do not know x^* , this conservatism is necessary under the framework of Batlle et al. (2023).

As we see in Section 3, our method addresses both of these challenges. In scenarios where evaluating Q_x is difficult, we sample a collection of design points in a bounded subset of the constraint set, sample the LLR under the null at each design point and use quantile regression to estimate the quantile surface. Furthermore, we can potentially remove these most conservative points from consideration by only considering parameter values that are not unlikely given the observed data.

3 Interval constructions

In this section, we present four related interval constructions to build on the theory developed in Batlle et al. (2023) by addressing the key aforementioned challenges. The first implementation challenge is to handle the potentially unbounded constraint set, for which we define and apply the "Berger–Boos" set to create a data-dependent subset of the original constraint set. This set leads to our four interval definitions, which follow a two-stage taxonomy; Global versus Sliced and Inverted versus Optimized. The second implementation challenge is computing these interval constructions in practice which we achieve using a combination of novel sampling algorithms and quantile regression. We will cover this in a later section (Section 5).

We rewrite the data-generating process articulated in the introduction:

$$\boldsymbol{y} = f(\boldsymbol{x}^*) + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}), \quad \boldsymbol{x}^* \in \mathcal{X}$$
 (19)

where $f : \mathbb{R}^p \to \mathbb{R}^n$ is a known forward model and $\operatorname{Cov}(\boldsymbol{\varepsilon}) = \boldsymbol{I}$, without loss of generality. Let $f^{-1}(A) := \{\boldsymbol{x} : f(\boldsymbol{x}) \in A\}$ be the pre-image of a set $A \subset \mathbb{R}^n$ under the forward map f.

Although the assumed data-generating process given by Equation (19) is less general than that of $\boldsymbol{y} \sim P_{\boldsymbol{x}^*}$ assumed in Section 2, it is still sufficiently general to contain the application areas mentioned in Section 1.1. Slightly generalizing the form of the additive noise distribution would make the construction of the Berger-Boos set more complicated, but is possible in principle.

3.1 Global and sliced confidence sets using the Berger–Boos set

Both $Q_{\mu,1-\alpha}^{\max}$ and $Q_{1-\alpha}^{\max}$ in Section 2 suffer from the same theoretical and practical concerns. Theoretically, they are both *conservative* in the sense that they must control type-1 error probability under the worst case for the truth (i.e., the parameter setting with the largest quantile). Practically, not only can Q_x be difficult to compute, but if the constraint set \mathcal{X} is unbounded, computing the aforementioned quantiles becomes even more difficult. These challenges both stem from the composite nature of the null hypothesis. Berger and Boos (1994) introduce a compelling solution in the context of hypothesis testing with nuisance parameters (a special case of hypothesis testing on $\varphi(\mathbf{x})$), which is to control type-1 error only over a data-informed region of the parameter space. Following their construction, we build a confidence interval that, instead of using the maximum $1 - \alpha$ quantile over \mathcal{X} in (14) (or $\mathcal{X} \cap \Phi_{\mu}$ in (13)), uses a larger quantile $(1 - \gamma \text{ with } \gamma < \alpha)$ but that is maximized over a smaller set. The construction follows a three-step process:

(1) Choose $\eta \in [0, \alpha]$ and build a $1 - \eta$ confidence interval for x^* , \mathcal{B}_{η} . Under the additive Gaussian noise assumption in Equation (19), letting $\Gamma_{\eta}(\boldsymbol{y}) := \{\boldsymbol{y}' \in \mathbb{R}^n : \|\boldsymbol{y} - \boldsymbol{y}'\|_2^2 \leq \chi_{n,\eta}^2\}$ we have

 $\mathbb{P}(f(\boldsymbol{x}^*) \in \Gamma_{\eta}(\boldsymbol{y})) = 1 - \eta$, hence

$$\mathcal{B}_{\eta} := f^{-1}(\Gamma_{\eta}(\boldsymbol{y})) \cap \mathcal{X} = \{ \boldsymbol{x} \in \mathcal{X} : \|\boldsymbol{y} - f(\boldsymbol{x})\|_{2}^{2} \le \chi_{n,\eta}^{2} \}$$
(20)

is a $1 - \eta$ confidence set for x^{*3} . We refer to this pre-image confidence set, \mathcal{B}_{η} , as the "Berger–Boos" set.

(2) Optimize the $1 - \gamma$ quantile of the test statistic only over \mathcal{B}_{η} (or $\mathcal{B}_{\eta} \cap \Phi_{\mu}$), instead of \mathcal{X} (or $\mathcal{X} \cap \Phi_{\mu}$). Here, $\gamma < \alpha$ is chosen to ensure calibration. As proved in lemma 3.1, $\gamma = \alpha - \eta$ is a valid choice.

(3) Use the obtained quantiles, which we define as

$$\bar{q}_{\gamma,\eta}(\mu) := \sup_{\boldsymbol{x} \in \mathcal{B}_{\eta} \cap \Phi_{\mu}} Q_{\boldsymbol{x}}(1-\gamma), \tag{21}$$

$$\bar{q}_{\gamma,\eta} := \sup_{\boldsymbol{x}\in\mathcal{B}_{\eta}} Q_{\boldsymbol{x}}(1-\gamma), \tag{22}$$

the construct the following sliced (sl) and global (gl) confidence sets:

$$C^{\rm sl}_{\alpha}(\boldsymbol{y}; \mathcal{B}_{\eta}) := \{ \mu : \lambda(\mu, \boldsymbol{y}) \le \bar{q}_{\gamma, \eta}(\mu) \},$$
(23)

$$C^{\rm gl}_{\alpha}(\boldsymbol{y}; \mathcal{B}_{\eta}) := \{ \mu : \lambda(\mu, \boldsymbol{y}) \le \bar{q}_{\gamma, \eta} \}.$$
(24)

Analogous to the presentation in Batlle et al. (2023), we define *sliced* and *global* max-quantiles that include the Berger–Boos set to control the type-1 error probability of Test (10). Both of these max-quantiles maximize a larger γ -quantile over the set of interest instead of the α -quantile as shown in (13) and (14).

The following Lemma (analogous to Lemma 2.1 in Batlle et al. (2023)) gives a sufficient condition to select the values γ and η to ensure a $1 - \alpha$ coverage.

Lemma 3.1 (Setting η and γ to guarantee $1 - \alpha$ coverage). Let $\alpha \in (0, 1)$ and define $\bar{q}_{\gamma,\eta}(\mu)$ according (21) and $C^{\rm sl}_{\alpha}(\boldsymbol{y}; \mathcal{B}_{\eta})$ according to (23). For $\eta \in (0, \alpha)$, $C^{\rm sl}_{\alpha}(\boldsymbol{y}; \mathcal{B}_{\eta})$ is a $1 - \alpha$ confidence set for $\mu^* = \varphi(\boldsymbol{x}^*)$ if $\gamma \leq \alpha - \eta$. The length of the obtained interval is a non-increasing function of γ so the tightest interval will be obtained when equality is satisfied.

Proof. See Appendix A.1.

The coverage guarantee implied by Lemma 3.1 also implies coverage for $C_{\alpha}^{\text{gl}}(\boldsymbol{y}; \mathcal{B}_{\eta})$, as shown in the following Corollary.

Corollary 3.2. In the setting of Lemma 3.1. $C_{\alpha}^{\text{gl}}(\boldsymbol{y}; \mathcal{B}_{\eta})$ is a $1 - \alpha$ confidence set for $\mu^* = \varphi(\boldsymbol{x}^*)$ if $\gamma \leq \alpha - \eta$.

Proof. See Appendix A.2.

Remark 1. Following the generality of the original construction, Lemma 3.1 generalizes to any $1 - \eta$ confidence set of x^* and test statistic calibrated with the $1 - \gamma$ quantile. In our particular

³Intersecting the pre-image $f^{-1}(\Gamma_{\eta}(\boldsymbol{y}))$ with the constraint set \mathcal{X} does not change the coverage probability since we know $\boldsymbol{x}^* \in \mathcal{X}$.

case, since both the confidence set \mathcal{B}_{η} and the test statistic $\lambda(\mu, y)$ depend on the Gaussian loglikelihood $\|\boldsymbol{y} - f(\boldsymbol{x})\|_2^2$, there is a stronger interpretation of the Berger-Boos construction. Whenever $\mathcal{B}_{\eta} \cap \Phi_{\mu} \neq \emptyset$, the Berger-Boos construction is equivalent to a data-dependant reduction of the constraint set, replacing \mathcal{X} for \mathcal{B}_{η} both in the test statistic and the quantile optimization problems. Since this is done after seeing the data, the optimized quantile needs to increase to $1 - \alpha + \eta$ to maintain $1 - \alpha$ coverage.

Navigating the Berger-Boos parameter choices and trade-offs. Setting $\eta = 0$ in $C_{\alpha}^{\text{gl}}(\boldsymbol{y}; \mathcal{B}_{\eta})$ we maximize $Q_{\boldsymbol{x}}(1-\gamma) = Q_{\boldsymbol{x}}(1-\alpha)$ over \mathcal{X} , which returns $Q_{1-\alpha}^{\max}$ from Equation (14). For $\eta > 0$, the Berger-Boos construction restricts the quantile maximization to a smaller subset of the parameter space while maximizing a larger $1-\gamma$ quantile to maintain the desired $1-\alpha$ confidence level of the final confidence set. Fundamentally, the choice of η reflects a trade-off between these two opposing effects: As we increase η from 0, the set over which the quantile function is optimized shrinks, but the quantile level is increased. In Section 6.2, we perform a numerical experiment comparing average interval length for different values of η , showing that a small $\eta > 0$ can be beneficial not only numerically (since in certain cases it allows sampling over a bounded set) but also in terms of the average length of the resulting interval.

3.2 Interval constructions

There are possibly many ways to compute $C_{\alpha}^{\text{gl}}(\boldsymbol{y}; \mathcal{B}_{\eta})$ and $C_{\alpha}^{\text{sl}}(\boldsymbol{y}; \mathcal{B}_{\eta})$ in practice. Obtaining either set comes down to the computation of $\bar{q}_{\gamma,\eta}(\mu)$ and $\bar{q}_{\gamma,\eta}$. If the quantile function $Q_{\boldsymbol{x}}$ and its gradient $\nabla_{\boldsymbol{x}} Q_{\boldsymbol{x}}$ could be evaluated, computing these quantities could potentially be achieved using a firstorder numerical optimizer. However, we emphasize that since the sliced max-quantile, $\bar{q}_{\gamma,\eta}(\mu)$, is a function of the level-set parameter, μ , such an optimization would have to be done for each possible functional value. Since such easy function and gradient evaluations rarely exist (see Batlle et al. (2023) for some examples where such evaluations are possible), this paper develops a *sampling*based approach to estimate these quantities to construct Confidence Sets (23) and (24). As we demonstrate below, once we estimate $\bar{q}_{\gamma,\eta}(\mu)$ and $\bar{q}_{\gamma,\eta}$, we can either use the output from the sampling algorithm to compute the Global or Sliced interval via classical test inversion, or we can use the estimated max-quantiles in optimizations similar to those in Interval (18). As such, we introduce four interval constructions; Global Inverted, Global Optimized, Sliced Inverted, and Sliced Optimized. These four options are summarized in Section 3.2.

We leverage our ability to sample $\lambda(\mu, \mathbf{y})$ by sampling $\mathbf{y} \sim \mathcal{N}(f(\mathbf{x}), \mathbf{I})$ to estimate the desired max-quantiles. We present two algorithms, both of which first generate a random set of design points from the Berger-Boos set. If we can directly *efficiently* compute the desired quantile at each design point, this capability is leveraged in Algorithm 1. If we cannot afford such a computation, Algorithm 2 presents an alternative, which first samples one realization of the LLR test statistic at each sampled design point, and then performs quantile regression with the generated pairs of design points and LLR values to estimate the underlying quantile surface. The sampled design points and either the exact or estimated quantiles at the design points are then used to produce the final intervals.

Both algorithms start by sampling \mathcal{B}_{η} uniformly at random to generate a collection of M design points across the Berger-Boos set, i.e., $\bar{\boldsymbol{x}}_1, \bar{\boldsymbol{x}}_2, \ldots, \bar{\boldsymbol{x}}_M \sim \mathcal{U}(\mathcal{B}_{\eta})$. Even under the additive Gaussian noise assumption where the pre-image $f^{-1}(\Gamma_{\eta}(\boldsymbol{y}))$ is an ellipsoid, this step is non-trivial since the ill-posedness can produce an extremely narrow or unbounded ellipsoid, resulting in sharp boundaries when intersected with the known constraints in addition to a large portion of the pre-image lying outside of \mathcal{X} , making rejection-sampling challenging. We address these sampling challenges in Section 5.1. If computing $\lambda(\mu, y)$ for a given μ and y is inexpensive, Algorithm 1 directly estimates the quantile function at each design point. This computation is most likely inexpensive when there are closed-form solutions to the LLR's subordinate optimizations. As shown in Algorithm 1, an easy way to estimate each design point's quantile is to sample its test statistic N times and take the appropriate percentile. More often, computing $\lambda(\mu, y)$ is expensive since it involves two constrained optimizations which are possibly non-convex due to either the constraints or the forward model, or just numerically challenging due to the ill-posedness of the problem. In this scenario, one can use Algorithm 2 to sample one realization of the test statistic at each design point and estimate the quantile function over the Berger–Boos set using quantile regression. While Algorithm 1 relies upon computational strength to compute the LLR $N \times M$ times, Algorithm 2 shifts complexity to the quantile regression and relieves the computational burden by assuming that there is information to be shared about the quantile surface between design points (namely that the quantile surface is smooth). The emphasis on the quantile regression further necessitates that the quantile regression be performed well. We discuss some considerations to this end in Section 5.2. We note that although Algorithm 1 and Algorithm 2 make use of the additive Gaussian noise assumption, they are not necessarily limited to this assumption given the proper adjustments to the Berger-Boos set, assuming one still has the ability to sample from the data-generating process. Since these approaches are sampling-based, it is necessary to show convergence as the number of samples gets large.

The approach of Algorithm 2 is inspired by recent uncertainty quantification approaches in likelihoodfree inference where one can sample from a likelihood but cannot easily compute it. Specifically, we draw inspiration from the use of quantile regression in Dalmasso et al. (2020, 2024); Masserano et al. (2023, 2024). Although these approaches differ in detail and implementation from our approach (e.g., they typically focus on settings with low-dimensional \boldsymbol{x}), they overlap in the sampling and quantile regression perspectives, which effectively allow machine learning to supplement for computationally intensive or intractable forward models. Namely, rather than assuming a purely stochastic forward model and therefore only being able to sample from the likelihood, we assume fis a deterministic function involved in $\lambda(\mu, \boldsymbol{y})$, which is a random quantity due to the additive noise. Since Algorithm 2 involves training a quantile regressor, it includes separate training and testing sets of design points over the Berger-Boos set and samples from their respective test statistics.

Algorithm 1 Direct estimation of quantiles

Input: $\alpha, \gamma, \eta \in (0, 1)$ such that $\gamma = \alpha - \eta, M, N \in \mathbb{N}$.

- 1: Construct Berger-Boos confidence set: Create $\Gamma_{\eta}(\boldsymbol{y}) \subseteq \mathbb{R}^{n}$ such that $\mathbb{P}(f(\boldsymbol{x}^{*}) \in \Gamma_{\eta}(\boldsymbol{y})) \geq 1 \eta$. $f^{-1}(\Gamma(\boldsymbol{y}))$ is also a 1η confidence set for \boldsymbol{x}^{*} , as is $\mathcal{B}_{\eta} = f^{-1}(\Gamma_{\eta}(\boldsymbol{y})) \cap \mathcal{X}$.
- 2: Sample from the Berger-Boos confidence set, \mathcal{B}_{η} : Sample $\bar{x}_1, \ldots, \bar{x}_M \sim \mathcal{U}(\mathcal{B}_{\eta})$.
- 3: for k = 1, 2, ..., M do
- 4: Sample noise realizations: Sample N noise realizations: $\varepsilon_1, \ldots, \varepsilon_N \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.
- 5: Sample from the LLR distribution: Create an ensemble of LLR draws under $\bar{\boldsymbol{x}}_k$: $\{\lambda_i\}_{i=1}^N$, where $\lambda_i := \lambda \left(\varphi(\bar{\boldsymbol{x}}_k), f(\bar{\boldsymbol{x}}_k) + \boldsymbol{\varepsilon}_i; \mathcal{X}\right)$.
- 6: **Compute percentile estimate of the** γ **-quantile**: Compute the $(1 \gamma) \times 100$ percentile of the LLR samples for the data generating process under $\bar{\boldsymbol{x}}_k$, i.e., $\hat{q}^k_{\gamma} := \lambda_{(\{(1-\gamma)N\})}$, where $\{\cdot\}$ denotes the nearest whole number and $\lambda_{(i)}$ denotes the *i*-th order statistic.

7: end for

Output: Pairs of sampled design points and their respective γ -quantiles, i.e., $\{(\bar{\boldsymbol{x}}_k, \hat{q}_{\gamma}^k)\}_{k=1}^M$.

Algorithm 2 Quantile regression estimate of quantile surface

Input: $\alpha, \gamma, \eta \in (0, 1)$ such that $\gamma = \alpha - \eta$; $M_{tr}, M \in \mathbb{N}$.

- 1: Construct Berger-Boos confidence set: Create $\Gamma_{\eta}(\boldsymbol{y}) \subseteq \mathbb{R}^{n}$ such that $\mathbb{P}(f(\boldsymbol{x}^{*}) \in \Gamma_{\eta}(\boldsymbol{y})) \geq 1 \eta$. $f^{-1}(\Gamma(\boldsymbol{y}))$ is also a 1η confidence set for \boldsymbol{x}^{*} , as is $\mathcal{B}_{\eta} = f^{-1}(\Gamma_{\eta}(\boldsymbol{y})) \cap \mathcal{X}$.
- 2: Sample from the Berger-Boos confidence set \mathcal{B}_{η} : Sample $\bar{x}_1, \ldots, \bar{x}_{M_{\text{tr}}} \sim \mathcal{U}(\mathcal{B}_{\eta})$ design points to train the quantile regressor and $\tilde{x}_1, \ldots, \tilde{x}_M \sim \mathcal{U}(\mathcal{B}_{\eta})$ test points to invert the interval, generating $M_{\text{tr}} + M$ total samples. Since the test points are used for the interval inversion, they are used as out-of-sample points for the quantile regressor.
- 3: for $k = 1, 2, ..., M_{tr}$ do
- 4: Sample a noise realization: Sample a noise realization: $\varepsilon_k \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.
- 5: Sample from the LLR distribution: Compute the LLR under \bar{x}_k with sampled noise ε_k : $\lambda_k := \lambda \left(\varphi(\bar{x}_k), f(\bar{x}_k) + \varepsilon_k; \mathcal{X} \right).$
- 6: **end for**
- 7: Estimate the quantile function using quantile regression: Using the generated pairs $\{(\bar{\boldsymbol{x}}_k, \lambda_k)\}_{k=1}^{M_{\text{tr}}}$, estimate the upper γ -conditional quantile function, $\hat{q}_{\gamma}(\boldsymbol{x})$, using quantile regression.

Output: Generate γ -quantile predictions at out-of-sample test points, $\{(\widetilde{\boldsymbol{x}}_k, \widehat{q}_{\gamma}(\widetilde{\boldsymbol{x}}_k))\}_{k=1}^M$.

With the generated pairs from Algorithms (1) and (2), we present two strategies to estimate each of $C^{\rm gl}_{\alpha}(\boldsymbol{y}; \mathcal{B}_{\eta})$ and $C^{\rm sl}_{\alpha}(\boldsymbol{y}; \mathcal{B}_{\eta})$. To streamline notation, let $(\boldsymbol{x}_k, q^k_{\gamma})$ denote the k-th pair from either algorithm. This is notationally helpful since Algorithm 1 only generates one set of parameter samples, whereas Algorithm 2 generates two. That is, using Algorithm 1, $\boldsymbol{x}_k := \bar{\boldsymbol{x}}_k$ and $q^k_{\gamma} := \hat{q}^k_{\gamma}$ and using Algorithm 2, $\boldsymbol{x}_k := \tilde{\boldsymbol{x}}_k$ and $q^k_{\gamma} := \hat{q}_{\gamma}(\tilde{\boldsymbol{x}}_k)$.

To estimate the global confidence set, we estimate $\bar{q}_{\gamma,\eta}$ using the empirical maximum, $\hat{q} := \max_k q_{\gamma}^k$. This results in the following two interval constructions:

$$C_{\text{opt}}^{\text{gl}}(\boldsymbol{y}) := \min / \max \left\{ \varphi(\boldsymbol{x}) : \boldsymbol{x} \in D(\widehat{q} + s(\boldsymbol{y})^2, \boldsymbol{y}) \right\}$$
(25)

$$C_{\text{inv}}^{\text{gl}}(\boldsymbol{y}) := \min / \max \left\{ \varphi(\boldsymbol{x}_k) : k = 1, \dots, M \text{ and } \lambda(\varphi(\boldsymbol{x}_k), \boldsymbol{y}; \mathcal{X}) \le \widehat{q} \right\},$$
(26)

We refer to Interval (26) as the Global Inverted interval construction since it is endpoints defined by only those sampled parameter values that comport with the maximum estimated quantile LLR cutoff. We refer to Interval (25) as the Global Optimized interval since its endpoints are defined by the extreme functional values of a feasible region defined by \hat{q} . Although, $C_{inv}^{gl}(\boldsymbol{y}) \neq C_{opt}^{gl}(\boldsymbol{y})$ due to finite sample, they are asymptotically equal and in practice show similar performance in terms of coverage and expected length, as shown in Section 6. We prove the consistency of the interval constructed via inversion and via optimization in Theorem 4.1. Since the endpoints of both the inverted and optimized intervals converge in probability to the endpoints of $C_{\alpha}^{gl}(\boldsymbol{y}; \mathcal{B}_{\eta})$, the two interval constructions are asymptotically equivalent.

To estimate the sliced set, we estimate $\bar{q}_{\gamma,\eta}(\mu)$ using a rolling maximum of the sampled q_{γ}^{k} values, as ordered by the sampled functional values, and directly accept or reject functional values based on their estimated quantile. These approaches result in two interval constructions:

$$C_{\text{opt}}^{\text{sl}}(\boldsymbol{y}) := \min / \max \left\{ \mu \in \mathbb{R} : \lambda(\varphi(\mu), \boldsymbol{y}; \mathcal{X}) \le \widehat{m}_{\gamma}(\mu) \right\},$$
(27)

$$C_{\text{inv}}^{\text{sl}}(\boldsymbol{y}) := \min / \max \left\{ \varphi(\boldsymbol{x}_k) : k = 1, \dots, M \text{ and } \lambda(\varphi(\boldsymbol{x}_k, \boldsymbol{y}; \mathcal{X}) \le q_{\gamma}^k) \right\},$$
(28)

Table 1: Summary of methods based on global and sliced max quantile, and whether they are optimizationbased or inversion-based.

Category	Global max quantile	Sliced max quantile
Inversion-based methods	Interval (26)	Interval (28)
Optimization-based methods	Interval (25)	Interval (27)

where $\widehat{m}_{\gamma}(\mu)$ denotes rolling estimate of $\overline{q}_{\gamma,\eta}(\mu)$ defined as follows. We refer to $I^{\rm sl}(\boldsymbol{y})$ as the "sliced" index set, i.e., those indices for which the LLR at a particular functional value is less than the estimated quantile at a design point generating that functional value.

The rolling maximum quantile is defined using estimated quantiles ordered by the sampled functional values. Choose "rolling" parameter, $T \in \mathbb{N}$, and let $\sigma(1), \sigma(2), \ldots, \sigma(M)$ define an ordering such that $\mu_{\sigma(k)} \leq \mu_{\sigma(k+1)}$ for all $k = 1, \ldots, M - 1$. Define $Q_k := \{q_{\gamma}^{\sigma(k)}, q_{\gamma}^{\sigma(k-1)}, \ldots, q_{\gamma}^{\sigma(k-T)}\}$ Then, for a given $\mu \in [\min_{\boldsymbol{x} \in \mathcal{B}_{\eta}} \varphi(\boldsymbol{x}), \max_{\boldsymbol{x} \in \mathcal{B}_{\eta}} \varphi(\boldsymbol{x})]$, define $k^*(\mu) := \operatorname{argmin}_k |\mu - \mu_{\sigma(k)}|$ and define $\widehat{m}_{\gamma}(\mu) := \max\{q \in Q_{k^*(\mu)}\}$. Using the rolling maximum quantile is one possible way to estimate $\bar{q}_{\gamma,\eta}(\mu)$. One could also bin the quantiles by functional value, compute the maximum predicted quantile in each bin, and then fit a nonparametric regression to fit the maximum binned quantiles to the functional values.

This estimator choice for $\hat{m}_{\gamma}(\mu)$ affects how one computes $C_{\text{opt}}^{\text{sl}}(\boldsymbol{y})$. To see why, note that we can re-express $C_{\text{opt}}^{\text{sl}}(\boldsymbol{y})$ as follows:

$$C_{\text{opt}}^{\text{sl}}(\boldsymbol{y}) := \min / \max \left\{ \varphi(\boldsymbol{x}) : \boldsymbol{x} \in D(\widehat{m}_{\gamma}(\varphi(\boldsymbol{x})) + s(\boldsymbol{y})^2, \boldsymbol{y}) \right\}.$$
(29)

As such, one could substitute the estimated $\hat{m}_{\gamma}(\cdot)$ into each endpoint optimization. However, this estimated curve is likely not convex in \boldsymbol{x} and therefore complicates the optimizations. One could also pursue a root-finding approach to find the set of μ such that $\lambda(\mu, \boldsymbol{y}) = \hat{m}_{\gamma}(\mu)$ since these intersection points define a set of accepted functional values. This approach can also be complex in proportion to the complexity of the estimated curve. In our view, the most pragmatic approach is to simply determine all μ_k such that $\lambda(\mu_k, \boldsymbol{y}; \mathcal{X}) \leq \hat{m}_{\gamma}(\mu_k)$, and then define the endpoints of $C_{\text{opt}}^{\text{sl}}(\boldsymbol{y})$ to be the minimum and maximum of those accepted sampled values. The consistency of the computed endpoints for both the inverted and optimized sliced intervals is proven in Theorem 4.1. Similar to the global interval constructions, the consistency of both constructions to $C_{\alpha}^{\text{sl}}(\boldsymbol{y}; \mathcal{B}_{\eta})$ also establishes the asymptotic equivalent of the inverted and optimized approaches.

Although both Algorithms 1 and 2 can produce a quantile estimate, it is worth noting a few key differences between the two. First, Algorithm 1 involves a nested sampling loop. Thus, any statistical guarantee regarding the validity of its output has the two moving parts of the accuracy of \hat{q}_{γ}^k as N gets large and the proximity of the approximated quantile to the true maximum quantile function value over \mathcal{B}_{η} as M gets large. By contrast, Algorithm 2 only has one sampling loop and estimates the full quantile function surface over \mathcal{B}_{η} , producing an estimate of the maximum as a consequence, but is reliant on smoothness to do an accurate regression with finitely many data points.

4 Theoretical justification

In this section, we provide convergence results for the interval constructions of the previous section. We prove that, under certain assumptions, whenever Algorithm 1 or Algorithm 2 are used to

estimate the maximum quantiles, the global interval constructions (25) and (26) converge in probability to the true global interval $C^{\rm gl}_{\alpha}(\boldsymbol{y}; \mathcal{B}_{\eta})$ in (24), and the sliced interval constructions (27) and (28) converge in probability to the true sliced interval $C^{\rm sl}_{\alpha}(\boldsymbol{y}; \boldsymbol{\mathcal{B}}_{\eta})$ in (23). Although our approach involving quantile regression is inspired by the approaches in Dalmasso et al. (2020, 2024); Masserano et al. (2023) and therefore requires similar theoretical results to connect the consistency of the quantile estimation with the interval validity, our approach is sufficiently different and requires novel theoretical insights. First, the inversion of a hypothesis test with a composite null adds a layer of complexity to the proofs. Dalmasso et al. (2024) does address composite null hypotheses, but in the context of nuisance parameters where they use a profile likelihood approach. Second and more fundamentally, Dalmasso et al. (2024); Masserano et al. (2023) construct confidence sets which are shown to achieve the desired coverage level asymptotically in the number of samples used to train the quantile regressor. By contrast, we show that our computed intervals converge in probability to the theoretical Intervals (23) and (24) which achieve the correct coverage by definition. Third and finally, since we use the sampled points to invert the interval, it is insufficient to show that the cutoff is consistent as the results in Dalmasso et al. (2020, 2024); Masserano et al. (2023) show. Our inverted intervals require proof that the samplers can get arbitrarily close to the true endpoint boundaries.

Throughout the following results, we assume that for fixed $\gamma \in (0, 1)$, $Q_x(1 - \gamma)$ is a continuous function of x. We refer the reader to Kibzun and Kan (1997) for a full analysis of the properties of parametrized quantile functions. We furthermore assume that \mathcal{B}_{η} is a compact set without isolated points, so that the sampling methods eventually sample close to every point.

Theorem 4.1. Assume that $Q_{\boldsymbol{x}}(1-\gamma)$ is a continuous function of \boldsymbol{x} , and let \mathcal{B}_{η} be compact and without isolated points. Let the quantile regression in Algorithm 2 be consistent for all \boldsymbol{x} , i.e, such that $\mathbb{P}(|\hat{q}_{\gamma}(\boldsymbol{x}) - Q_{\boldsymbol{x}}(1-\gamma)| > \varepsilon) \to 0$ as $M_{tr} \to \infty$ is satisfied $\forall \varepsilon > 0$. We will write $\stackrel{\mathrm{P}}{\to}$ for convergence in probability, understood as $N, M \to \infty$ if Algorithm 1 is used and as $M_{tr}, M \to \infty$ if Algorithm 2 is used. For either algorithm, we have, for a given observation \boldsymbol{y} :

- 1. $C_{\text{inv}}^{\text{gl}}(\boldsymbol{y}) \xrightarrow{\text{p}} C_{\alpha}^{\text{gl}}(\boldsymbol{y}),$
- 2. $C_{\text{inv}}^{\text{sl}}(\boldsymbol{y}) \xrightarrow{\mathrm{p}} C_{\alpha}^{\text{sl}}(\boldsymbol{y}).$

Further assume that there exists a point $\bar{\mu} \in \varphi(\mathcal{B}_{\eta})$ satisfying $\lambda(\bar{\mu}, \boldsymbol{y}; \mathcal{X}) < \bar{q}_{\gamma, \eta}$, then

3. $C_{\text{opt}}^{\text{gl}}(\boldsymbol{y}) \xrightarrow{\text{p}} C_{\alpha}^{\text{gl}}(\boldsymbol{y}).$

Finally, further assuming technical conditions discussed in Appendix B, then

4. $C^{\mathrm{sl}}_{\mathrm{opt}}(\boldsymbol{y}) \xrightarrow{\mathrm{p}} C^{\mathrm{sl}}_{\alpha}(\boldsymbol{y}).$

Proof. See Appendix B.

Note that if $\lambda(\mu, \boldsymbol{y})$ is convex in μ , as it is for linear forward models and quantities of interest (Batlle et al., 2023, Proposition 2.5), the condition of the existence of $\bar{\mu}$ satisfying $\lambda(\bar{\mu}, \boldsymbol{y}; \mathcal{X}) < \bar{q}_{\gamma,\eta}$ is equivalent to the interval not being empty. For nonlinear forward models, it is a slightly stronger condition. For the convergence of the sliced optimized version, one needs to show that $\inf_{\mu:\lambda(\mu,\boldsymbol{y})\leq \widehat{m}_{\gamma}(\mu)}\mu$ converges to $\inf_{\mu:\lambda(\mu,\boldsymbol{y})\leq m_{\gamma}(\mu)}\mu$ as $\widehat{m}_{\gamma}(\mu)$ converges to $m_{\gamma}(\mu)$. In order to do so, we study the convergence of optimization problems of the form $\inf_{\mu:\widehat{f}(\mu)\geq 0}\mu$ to $\inf_{\mu:f(\mu)\geq 0}\mu$ as \widehat{f} converges to f. Although the result is not true in general, we provide sufficient technical conditions

about f and the uniformness of convergence of \hat{f} to f for the result to hold. We discuss the details in Appendix B.

5 Implementation methodology

5.1 Sampling the pre-image Berger–Boos set

The viability of this method directly relies upon our ability to sample from the Berger-Boos set, $\mathcal{B}_{\eta} = f^{-1}(\Gamma_{\eta}(\boldsymbol{y})) \cap \mathcal{X}$. Under the assumed data generating process in Equation (19), the Berger-Boos set is defined as follows:

$$\mathcal{B}_{\eta} = f^{-1}(\Gamma_{\eta}(\boldsymbol{y})) \cap \mathcal{X} = \left\{ \boldsymbol{x} \in \mathcal{X} : (\boldsymbol{y} - f(\boldsymbol{x}))^{\top} \boldsymbol{\Sigma}^{-1}(\boldsymbol{y} - f(\boldsymbol{x})) \le \chi_{n,\eta}^{2} \right\},$$
(30)

where we have generalized to a non-identity covariance matrix, Σ , to make the following exposition more general. This set is equivalent to the set over which the strict bounds intervals are optimized in Stark (1992b), also called "SSB" intervals in Stanley et al. (2022); Batlle et al. (2023). Note, however, that one would use $\chi^2_{n,\alpha}$ instead of $\chi^2_{n,\eta}$ for $1 - \alpha$ interval computation in that scenario.

We discuss sampling \mathcal{B}_{η} in the linear-Gaussian case, where $f(\mathbf{x}) = \mathbf{K}\mathbf{x}$, as this scenario aligns with the numerical experiments presented in Section 6. We present two sampling approaches; the "Voelker-Gossman-Stewart" (VGS) algorithm based on Voelker et al. (2017) for efficiently sampling ellipsoids uniformly at random in low-dimensional scenarios and an MCMC-based algorithm we call the Polytope sampler based on the Vaidya walk presented in Chen et al. (2018). When f is linear and has full column rank, \mathcal{B}_{η} is an ellipsoid intersected with the constraint set, making the VGS algorithm an effective option in low dimensional settings (see Section 5.1.1). In all other scenarios, especially higher dimensional ones, the Polytope sampler is a better option, as naive accept-reject algorithms become intractable (see Section 5.1.2).

5.1.1 VGS Sampler for low-dimensional and full column rank settings

Under the linear-Gaussian assumptions, when K has full column rank, it can be shown that

$$\mathcal{B}_{\eta} = \left\{ \boldsymbol{x} \in \mathcal{X} : (\boldsymbol{x} - \widehat{\boldsymbol{x}})^{\top} \boldsymbol{K}^{\top} \boldsymbol{\Sigma}^{-1} \boldsymbol{K} (\boldsymbol{x} - \widehat{\boldsymbol{x}}) \le \chi_{n,\eta}^{2} \right\},$$
(31)

where \hat{x} is the Generalized Least-Squares (GLS) estimator. Equation (31) describes an ellipsoid in \mathbb{R}^p intersected with \mathcal{X} with axis directions and lengths determined by $\chi^2_{n,\eta}$ and the eigenvectors and eigenvalues of $\mathbf{K}^{\top} \boldsymbol{\Sigma}^{-1} \mathbf{K}$, respectively (see Rust and Burrus (1972) for more details). An ellipsoid is nothing but a deformed ball. As such, we can sample uniformly at random from the ellipsoid in \mathbb{R}^p using an algorithm sampling uniformly at random from a p-ball, followed by an appropriate linear transformation and translation. We can then include an additional accept-reject step to account for the constraint set \mathcal{X} . Note, all subsequent discussions of spheres and balls assume unit radii and centering at the origin. Also note that because the ellipsoid in (31) is centered around the GLS estimator, we can sample points first from within an ellipsoid of the same shape centered at the origin, and then translate those points by \hat{x} .

Voelker et al. (2017) propose a particularly efficient and clever algorithm to sample uniformly at random from the *p*-ball by proving a connection between uniform sampling on the (p + 1)-sphere (i.e., the surface of the ball in \mathbb{R}^{p+2}) and uniform sampling within the *p*-ball (i.e., within the unit ball in \mathbb{R}^p). Namely, one can sample a Gaussian $\tilde{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{p+2})$ and normalize $\mathbf{z} := \tilde{z}/\|\tilde{z}\|_2$ to sample a point uniformly at random from the (p + 1)-sphere. Next, one simply drops the last two

elements of z to obtain a point that is uniformly sampled from within a p-ball. The validity of this approach is substantiated by Lemma 1 and Theorem 1 in Voelker et al. (2017) and relies upon a distribution result about the ratio of chi-squared distributions and preservation of distribution under an orthogonal transformation. To denote sampling a point following this procedure, we use the notation $x \sim \text{VGS}(p)$.

To sample from our desired ellipsoid in (31), first consider the eigendecomposition of $\mathbf{K}^{\top} \mathbf{\Sigma}^{-1} \mathbf{K} = \mathbf{P} \mathbf{\Omega} \mathbf{P}^{\top}$, where $\mathbf{\Omega} = \text{diag} \left(\omega_1^2, \omega_2^2, \dots, \omega_p^2 \right)$, ω_i is the *i*-th eigenvalue of $\mathbf{K}^{\top} \mathbf{\Sigma}^{-1} \mathbf{K}$ and \mathbf{P} is an orthonormal matrix where the columns vectors are the eigenvectors of $\mathbf{K}^{\top} \mathbf{\Sigma}^{-1} \mathbf{K}$. Denote $\mathbf{\Omega}^{1/2} = \text{diag} \left(\omega_1, \omega_2, \dots, \omega_p \right)$ and by extension, $\mathbf{\Omega}^{-1/2} = \text{diag} \left(\omega_1^{-1}, \omega_2^{-1}, \dots, \omega_p^{-1} \right)$ when $\omega_i > 0$ for all *i*. These decompositions imply the correct transformation to apply to points sampled from the *p*-ball via $\mathbf{x} \sim \text{VGS}(p)$. Namely, define $\mathbf{w} := \sqrt{\chi_{n,\eta}^2} \mathbf{P} \mathbf{\Omega}^{-1/2} \mathbf{x}$. To know that \mathbf{w} is sampled from the correct ellipsoid, it should be the case that $\mathbf{w}^{\top} \mathbf{K}^{\top} \mathbf{\Sigma}^{-1} \mathbf{K} \mathbf{w} \leq \chi_{n,\eta}^2$, as then we can simply make the update $\mathbf{w} \leftarrow \mathbf{w} + \hat{\mathbf{x}}$ to ensure that we have a point sampled from (31). This guarantee is verified as follows:

$$\boldsymbol{w}^{\top}\boldsymbol{K}^{\top}\boldsymbol{\Sigma}^{-1}\boldsymbol{K}\boldsymbol{w} = \boldsymbol{w}^{\top}\boldsymbol{P}\boldsymbol{\Omega}^{1/2}\boldsymbol{\Omega}^{1/2}\boldsymbol{P}^{\top}\boldsymbol{w}$$

$$= \chi_{n,\eta}^{2} \cdot \boldsymbol{x}^{\top}\boldsymbol{\Omega}^{-1/2}\boldsymbol{P}^{\top}\boldsymbol{P}\boldsymbol{\Omega}^{1/2}\boldsymbol{\Omega}^{1/2}\boldsymbol{P}^{\top}\boldsymbol{P}\boldsymbol{\Omega}^{-1/2}\boldsymbol{x}$$

$$= \chi_{n,\eta}^{2} \cdot \boldsymbol{x}^{\top}\boldsymbol{x} \leq \chi_{n,\eta}^{2},$$

(32)

where $\mathbf{P}^{\top}\mathbf{P} = \mathbf{I}$ by definition and the last line follows since we know \mathbf{x} is sampled from the p-ball and therefore $\mathbf{x}^{\top}\mathbf{x} \leq 1$. Equation (32) justifies that the \mathbf{w} samples lie within the desired ellipsoid, and their uniform distribution follows from the uniform distribution of \mathbf{x} in the p-ball since the distribution is preserved under a linear transformation. Finally, we accept \mathbf{w} if $\mathbf{w} \in \mathcal{X}$ and reject if $\mathbf{w} \notin \mathcal{X}$. This procedure for sampling at random from \mathcal{B}_{η} is summarized in Algorithm 3.

Algorithm 3 VGS Sampler

 $\underline{\text{Input}}: \ p \in \mathbb{N}, \ M \in \mathbb{N}, \ \boldsymbol{P}, \ \boldsymbol{\Omega}^{-1/2}, \ \widehat{\boldsymbol{x}}, \ \chi^2_{n,\eta}.$

1: Define $\mathcal{S} := \{\}$ to be the initialized set in which the sampled points are to be placed.

- 2: for k = 1, 2, ..., M do
- 3: Sample from the *p*-ball using VGS: $x_k := z_{1:p}/||z||_2$, where $z \sim \mathcal{N}(0, I_{p+2})$ and $z_{1:p}$ denotes taking the first through *p*-th indices (inclusive).
- 4: Transform VGS output: $\boldsymbol{w}_k := \sqrt{\chi_{n,\eta}^2 \boldsymbol{P} \boldsymbol{\Omega}^{-1/2} \boldsymbol{x}_k}.$
- 5: Translate w_k by the GLS estimator: $w_k \leftarrow w_k + \widehat{x}$.
- 6: Accept-Reject to incorporate constraints: If $w_k \in \mathcal{X}$, then add $\mathcal{S} \leftarrow \mathcal{S} \cup \{w_k\}$, else start loop iteration k again.

7: end for

Output: S containing uniformly sampled points over \mathcal{B}_{η} (as defined in (31)).

Generating samples using the VGS Sampler is efficient, but its feasibility diminishes in high dimensions because of the accept-reject step. To illustrate this point, consider the simple scenario where $\mathcal{X} = \mathbb{R}_{+}^{p}$, i.e., the non-negative orthant of \mathbb{R}^{p} and suppose we sample from the *p*-ball intersected with \mathbb{R}_{+}^{p} by sampling $\boldsymbol{x} \sim \text{VGS}(p)$ where $\boldsymbol{P} = \boldsymbol{\Omega} = \boldsymbol{I}$, $\hat{\boldsymbol{x}} = \boldsymbol{0}$ and we use 1 instead of $\chi_{n,\eta}^{2}$. Then, $\mathbb{P}(\boldsymbol{x} \in \mathbb{R}_{+}^{p}) = 2^{-p}$ and the acceptance probability of each sample goes to zero exponentially in *p*. To make this point slightly more general, we generate data $\boldsymbol{y} \sim \mathcal{N}(\boldsymbol{x}_{p}^{*}, \boldsymbol{I}_{p})$ where $\boldsymbol{x}_{p}^{*} \in \mathbb{R}_{+}^{p}$ and is a vector of ones. For a collection of dimensions $p \in [2, 30]$, we sample $\boldsymbol{x} \sim \text{VGS}(p)$ with $\boldsymbol{P} = \boldsymbol{\Omega} = \boldsymbol{I}$ and $\hat{\boldsymbol{x}} = \boldsymbol{y}$ to estimate the probability that \boldsymbol{x} is in \mathbb{R}_{+}^{p} and plot the results in the left panel of Figure 5.1. Since the acceptance probability decays exponentially (under 10^{-4} for 30 dimensions), it becomes clear that this algorithm is inefficient in dimensions larger than 10, since the acceptance probability quickly becomes prohibitively small in regimes where a larger sample size is even more important.

Dimension is only one of two primary complicating factors. Not only is less of a (potentially) shifted p-ball's volume in the non-negative orthant as p grows but if our data are generated via $y \sim \mathcal{N}(Kx^*, I_p)$ with $x^* \in \mathbb{R}^p_+$ where the condition number of K is large, it is possible that even less of the ellipsoid from which the VGS Sampler draws points intersects with the parameter constraint. To illustrate this point, we generate a single observation from the aforementioned model using a $K \in \mathbb{R}^{40 \times 40}$ as described in Section 6.3. The $x^* \in \mathbb{R}^{40}_+$ is created in the same way as that of Section 6.3.1. The right panel of Figure 5.1 shows a computed probability mass function for the number of coordinates in a VGS Sampler draw (with P and Ω defined such that $K^{\top}K = P\Omega P^{\top}$) complying with the non-negativity parameter constraints. Equivalently, the right panel of Figure 5.1 shows the computed probability mass function for the number of coordinates lying within the Berger-Boos set. For this particular setup, we critically note that out of 5×10^4 draws from the VGS Sampler, none of the draws had all coordinates comply with the non-negativity constraint. By contrast, for the aforementioned noise model where we more simply sample from the p-ball intersected with the non-negative orthant, the computed acceptance probability is approximately 3.35×10^{-6} , in both cases emphasizing the VGS Sampler's poor performance in high-dimensional and ill-conditioned forward model regimes.



Figure 5.1: Numerical illustrations of the VGS Sampler's infeasibility in high dimensional regimes. The **left** panel shows the computed acceptance probability of a point drawn by the VGS Sampler with data generated from a non-negatively constrained Gaussian noise model. Crucially, at only 30 dimensions, the acceptance probability is already less than 10^{-4} for this particular setup. The **right** panel shows the computed probability mass function for the number of non-negative constraint complying coordinates of a VGS sample with data generated from a non-negatively constrained linear Gaussian model in 40 dimensions with a non-identity forward model. Since this is an example using a forward model with a large condition number ($\approx 1.6 \times 10^4$), we critically note that there is empirically zero probability of generating a sample within the non-negativity constraints.

5.1.2 Polytope sampler for general settings

In settings where the forward model is not linear and full column rank, Algorithm 3 fails. The ineffectiveness of this algorithm expands if the condition number of the linear forward model is large such that most of the pre-image ellipsoid defining the Berger–Boos set lies outside of the constraint set. Although these scenarios induce particular geometric challenges, in the linear-Gaussian case with convex \mathcal{X} we are still fundamentally sampling a convex set for which there is a vast

literature. For example, there is a vast sampling literature for computing Bayesian posteriors in high dimensions via nested sampling (Skilling, 2004; Ashton et al., 2022; Buchner, 2023). In particular, nested sampling has been successfully applied in high-dimensional cosmology settings using sophisticated approaches to strategically sample the parameter space by restricting prior sampling in various ways (Buchner, 2023; Montel et al., 2023). Although these approaches provide tools addressing a sampling setting similar to ours (i.e., the Berger–Boos set can be viewed as a portion of the parameter space defined by a cutoff on the likelihood) they are ultimately aimed at sampling from a particular distribution (i.e., the posterior), which is a stronger criterion than required here. For sampling general convex sets, simple algorithms like Hit-and-Run are available (Smith, 1984; Lovasz, 1999; Lovasz and Vempala, 2006). However, in the particular case considered here, more sophisticated and efficient algorithms can be devised. In particular, there exists a deep literature on random walks over polytopes such that the asymptotic stationary distribution of the walk is a uniform distribution over the polytope of interest (Kannan and Narayanan, 2012; Narayanan, 2016; Chen et al., 2018). As such, we propose to first construct a bounding polytope, \mathcal{P}^d composed of d half-spaces around \mathcal{B}_{η} , sample C random walks within the Berger–Boos set using the Vaidya walk as described in Chen et al. (2018) each starting at a point from a collection of strategically chosen locations, and then combine the parallel chains to create the final sample set. The detailed algorithm can be seen in Algorithm 4.

Although MCMC algorithms are typically evaluated using trace plots on individual dimensions of the parameter space, given the high-dimensionality of the problem and the final step combining several MCMC chains in the parameter space started from different positions, it is more meaningful to evaluate this algorithm's ability to sample fully from the functional space. Namely, we can solve for the largest and smallest values the functional can take within the Berger–Boos set as follows:

$$I_{\mathsf{BB}}(\boldsymbol{y}) := \left[\mu_{\mathsf{BB}}^{l}, \mu_{\mathsf{BB}}^{u}\right] = \left[\min_{\boldsymbol{x}\in\mathcal{B}_{\eta}}\varphi(\boldsymbol{x}), \max_{\boldsymbol{x}\in\mathcal{B}_{\eta}}\varphi(\boldsymbol{x})\right].$$
(33)

When $\varphi(\boldsymbol{x}) = \boldsymbol{h}^{\top} \boldsymbol{x}$ for some $\boldsymbol{h} \in \mathbb{R}^{p}$, $I_{\mathsf{BB}}(\boldsymbol{y})$ corresponds to the SSB interval in Stanley et al. (2022). Computing the sets $C_{\alpha}^{\mathrm{sl}}(\boldsymbol{y}; \mathcal{B}_{\eta})$ and $C_{\alpha}^{\mathrm{gl}}(\boldsymbol{y}; \mathcal{B}_{\eta})$ well is then contingent upon sampling functional values within $I_{\mathsf{BB}}(\boldsymbol{y})$ well since a functional value can only be included in the inverted set if the sampler has a non-zero probability of sampling arbitrarily close to it. By "well", we informally mean that the sampled functional values range at least between the endpoints of $I_{\mathsf{BB}}(\boldsymbol{y})$ and that if we partition $I_{\mathsf{BB}}(\boldsymbol{y})$ into n sub–intervals, that most if not all of the sub–intervals contain at least one sample. In practice, it will often be the case that the sampled functional values can lie outside the endpoints of $I_{\mathsf{BB}}(\boldsymbol{y})$ since the MCMC chains are sampling a bounding polytope of the Berger–Boos set.

Choosing the bounding polytope. Any bounded polytope of \mathcal{B}_{η} is the intersection of a finite number of half-spaces defined in \mathbb{R}^p . The task of constructing a bounding polytope is then equivalent to choosing a collection of d hyperplanes in \mathbb{R}^p to construct a set $\mathcal{P}^d := \{ \boldsymbol{x} \in \mathbb{R}^p : \boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{b} \}$ such that $\mathcal{B}_{\eta} \subseteq \mathcal{P}^d$, where $\boldsymbol{A} \in \mathbb{R}^{d \times p}$ and $\boldsymbol{b} \in \mathbb{R}^d$. Let $\boldsymbol{a}_i^{\top} \in \mathbb{R}^p$ denote the *i*-th row vector of \boldsymbol{A} . We compute b_i (the *i*-th element of \boldsymbol{b}) as

$$b_i = \max_{\boldsymbol{x} \in \mathcal{B}_\eta} \boldsymbol{a}_i^\top \boldsymbol{x}.$$
 (34)

This construction ensures the necessary inclusion. We consider three approaches to pick the vectors a_i : (i) using the constraints \mathcal{X} , (ii) using the known eigenvectors defining the bounded directions of the pre-image ellipsoid, and (iii) randomly. In practice, we combine these approaches to ensure that we consider only parameter settings in agreement with our physical constraints, and to tighten

the bounding Berger–Boos set polytope as much as possible. There is a tradeoff with respect to the latter consideration since the mixing time and computational cost for the Vaidya walk increase with the number of hyperplanes (Chen et al., 2018).

To incorporate the known parameter constraints, consider the non-negativity constraint used in Section 6, i.e., $\boldsymbol{x} \in \mathbb{R}^p_+$. To enforce non-negativity, we set $\boldsymbol{a}_i = -\boldsymbol{e}_i$ for $i = 1, \ldots, p$, where \boldsymbol{e}_i is defined by its *i*-th element set to one and the rest of its elements set to zero and $b_i = 0$. These choices produce p rows in A corresponding to the desired lower bounds (i.e., $x_i \ge 0$ for all i), but we can compute an additional p constraints using (34) with $a_i := e_i$ for $i = p + 1, \dots 2p$. These 2p constraints define a hyperrectangle enclosing the Berger-Boos set in the parameter space. To incorporate polytope constraints based upon the forward model, we use the ellipsoidal definition of the pre-image as shown in (31) and eigendecomposition of $\mathbf{K}^{\top} \boldsymbol{\Sigma}^{-1} \mathbf{K} = \mathbf{P} \boldsymbol{\Omega} \mathbf{P}^{\top}$ shown in (32). Note that this ellipsoid form is valid under the linear-Gaussian noise assumption and for all K. In the event that K is not full column rank, the ellipsoid defined by Equation (31) is still defined via the pseudo-inverse of of $K^{\top}\Sigma K$ but where the ellipsoid is unbounded in some directions. The column vectors of P corresponding to the non-zero entries on the diagonal of Ω are the eigenvectors corresponding to the bounded principal axes. As such, both p_i and $-p_i$ for $i = 1, \ldots, p$ (the column vectors of P) can be used as rows of A with their corresponding bounds defined by (34). Finally, to further tighten the polytope around the Berger-Boos set, we sample a multivariate Gaussian, i.e., $a_i \sim \mathcal{N}(\mathbf{0}, I_p)$ to include random hyperplanes. Note, it is possible that the unbounded directions of the ellipsoid defined by Equation (31) are not bounded when intersecting with the parameter constraint set and hence there is no bounding polytope. In this case, our interval construction should be unbounded since there is not enough information to produce a bounded confidence set for the quantity of interest.

Once the components (\mathbf{A}, \mathbf{b}) have been defined using some or all of the above hyperplane generation strategies, the Vaidya walk can immediately be employed to perform a random walk around the polytope. The primary intuitive requirement we wish to satisfy with any sampling scheme in this context is that every region of the Berger-Boos set has a non-zero probability of being sampled. Asymptotically, the Vaidya walk samples the desired polytope uniformly at random, which satisfies a stronger requirement. In practice, the need to sample non-uniformly can arise if there are particularly meaningful parameter settings in regions that are difficult for the random walk to reach. We explore such a case in Section 6.2. Additionally, although the asymptotic distribution of the Vaidya sampler is theoretically sufficient, in practice it often has some difficulty reaching the corners of the generated polytope. Although MCMC chain mixing is typically evaluated by looking at trace plots, this diagnostic is insufficient here because the dimension is high and the defined polytope can make different dimensions difficult to compare. Instead, we consider how well the sampler samples the functional values $\varphi(\mathbf{x})$. This motivates taking a collection of starting points between the SSB endpoints and the Chebyshev center of the polytope, as described in the following section.

Constructing the parallel chain starting points. Although the random walks defined in Chen et al. (2018) asymptotically sample uniformly over \mathcal{P}^d , given the long and thin shape of the preimage, running the Vaidya walk from even a "good" starting position does not consistently sample the functional space well. Instead, we use the following heuristic to construct several parallel chains that constitute a complete sample when combined. We define an even number of starting points, $C \in$ $2\mathbb{N}$, to roughly span the Berger-Boos set, run the Vaidya walk for M_p steps from each starting point to collect parameter settings $\{\tilde{x}_1, \ldots, \tilde{x}_{M_p}\}$, and combine the samples from each walk, resulting in $C \times M_p$ total samples. In practice, C and M_p are chosen to yield a total of M samples as desired for Algorithm 1 or Algorithm 2. Since we have designed this process to avoid the neccessity of any individual chain reaching its asymptotic distribution, we do not need a burn-in period for any chain as long as we are confident that the union of the sampled points sufficiently covers Interval (33) so that we may accept or reject any quantity of interest value in that interval. We define the collection of starting points using the line segments defined by the parameters generating the endpoints of $I_{\mathsf{BB}}(\boldsymbol{y})$ and the Chebyshev center of \mathcal{P}^d as defined in Boyd and Vandenberghe (2004), and denoted by \boldsymbol{x}_c . This point is the center of the largest ball contained within \mathcal{P}^d and therefore acts as a way to characterize the center of \mathcal{P}^d . We denote the parameter settings generating the endpoints of $I_{\mathsf{BB}}(\boldsymbol{y})$ by $\hat{\boldsymbol{x}}^l$ for the lower endpoint and $\hat{\boldsymbol{x}}^u$ for the upper endpoint. We then define a uniform grid of values $\{\tau_l\}_{l=1}^{C/2}$ such that $\tau_l \in (0, 1)$ and $\tau_l < \tau_{l+1}$ for all l. For each $k \in [C]$, define $k' := \lceil k/2 \rceil$ and set $\boldsymbol{x}_k^{start} := \tau_{k'} \hat{\boldsymbol{x}}^l + (1 - \tau_{k'}) \boldsymbol{x}_c$ if k is odd and $\boldsymbol{x}_k^{start} := \tau_{k'} \hat{\boldsymbol{x}}^u + (1 - \tau_{k'}) \boldsymbol{x}_c$ if k is even. Creating starting positions along the lines connecting these endpoints and the Chebyshev center accommodates the chosen polytope while helping ensure that samples are chosen spanning the range of possible functional values over the Berger–Boos set.

The empirical performance of the Polytope sampler can be seen in Figure 5.2, showing (left) a histogram of the functional values sampled and (right) a trace plot for the sampled Vaidya walks starting from points constructed as described above. These plots are generated for one observation from the 80-dimensional ill-posed inverse problem in the coverage study performed in Section 6.3. Critically, the histogram shows that the Polytope sampler samples the functional space well and the trace plot shows that our starting point construction heuristic performs well in practice.



Figure 5.2: Polytope sampler output for a realization of the 80-dimensional ill-posed inverse problem studied in Section 6.3. The **left** panel contains a histogram of sampled functional values which both span and cover well the range of $I_{BB}(\boldsymbol{y})$ (shown by the dashed gray lines in both plots). The **right** panel contains trace plots of the 14 Vaidya walks (each indicated by a different color) which together constitute the full sample. Our heuristic for choosing starting points along the lines connecting the parameter settings generating the endpoints of $I_{BB}(\boldsymbol{y})$ and the Chebyshev center of the polytope provides a good initial spread of starting functional values.

Sampling from the Berger-Boos set. With \mathcal{P}^d and the starting positions defined, we construct a sample $\mathcal{S} := \{\tilde{x}_1, \ldots, \tilde{x}_{C \times M_p}\}$. We leave the details of the Vaidya walk to the original paper Chen et al. (2018), but note an essential radius tuning parameter of the algorithm that must be chosen. This radius impacts the spread of a Gaussian proposal distribution for the walk and thus affects a new proposed point's acceptance probability. Choosing a too large radius results in a low acceptance rate of proposed steps, thus creating walk that does not mix well. In contrast, choosing a too small radius results in a high acceptance rate with relatively small step sizes. In practice, we find a radius setting of 0.5 works well as it produces an acceptance probability of $\approx 33.3\%$, producing a reasonable trade off between taking meaningful steps and not rejecting too many steps.

Algorithm 4 Polytope sampler

 $\underline{\text{Input}}: \ M_p \in \mathbb{N} \text{ and } C \in 2\mathbb{N}. \ \boldsymbol{A} \in \mathbb{R}^{d \times p}, \ \boldsymbol{b} \in \mathbb{R}^d. \ \{\tau_l\}_{l=1}^{C/2}, \text{ where } \tau_l \in (0,1) \text{ and } \tau_l < \tau_{l+1} \text{ for all } l.$

1: Let $S := \{\}$ be the set in which we store all sampled parameter settings.

2: Construct Chebyshev center: Solve for x_c using the following optimization.

$$\begin{array}{ll} \underset{\boldsymbol{x}_{c},r}{\operatorname{maximize}} & r\\ \text{subject to} & \boldsymbol{a}_{i}^{\top}\boldsymbol{x}_{c}+r\|\boldsymbol{a}_{i}\|_{2}\leq b_{i}, \ i=1,\ldots,d. \end{array}$$

3: Compute φ extremes of Berger-Boos set: Extreme points are computed with respect to the functional of interest:

$$\widehat{\boldsymbol{x}}^{l} := \operatorname{argmin} \ \varphi(\boldsymbol{x}) \quad \text{subject to} \ \boldsymbol{x} \in \mathcal{B}_{\eta},$$
$$\widehat{\boldsymbol{x}}^{u} := \operatorname{argmax} \ \varphi(\boldsymbol{x}) \quad \text{subject to} \ \boldsymbol{x} \in \mathcal{B}_{\eta}.$$
(35)

- 4: for k = 1, 2, ..., C do
- 5: Construct starting point: Define $k' := \lceil k/2 \rceil$. Define $\boldsymbol{x}_k^{start} = \tau_{k'} \hat{\boldsymbol{x}}^l + (1 \tau_{k'}) \boldsymbol{x}_c$ if k is odd and $\boldsymbol{x}_k^{start} = \tau_{k'} \hat{\boldsymbol{x}}^u + (1 \tau_{k'}) \boldsymbol{x}_c$ if k is even.
- 6: Run the Vaidya walk for M_p steps: Collect samples $\{\tilde{x}_1, \ldots, \tilde{x}_{M_p}\}$ and add them to S. 7: end for

Output: S containing sampled points over \mathcal{B}_{η} .

5.2 Quantile regression

Algorithm 2 explained in Section 3.2 involves using quantile regression to learn a quantile surface from a collection of pairs of design points and samples from the LLR test statistic. As previously mentioned, similar approaches have been taken in Dalmasso et al. (2020, 2024); Masserano et al. (2023, 2024), and since quantile regression is a technique facilitating our interval constructions, we will only give a brief overview of quantile regression and some different ways to implement it. Fundamentally, given a one-dimension random variable $z \sim P_x$ that depends on the parameter $x \in \mathbb{R}^p$, we are interested in the upper γ -quantile at every parameter setting, i.e.,

$$\mathbb{P}_{\boldsymbol{x}}(z > Q_{\boldsymbol{x}}(1-\gamma)) = \gamma. \tag{36}$$

We note that the quantile surface itself is not random, so we can use draws from the distribution P_x to estimate Q_x at a given parameter setting x. However, as noted in Section 3.2, performing such an estimate is not always computationally feasible, and intuitively, we might expect quantiles to vary smoothly over the parameter space, which would imply that information about a quantile at x_1 should be related to a quantile at x_2 if these points are close. In statistics literature, estimating Q_x is framed as estimating a quantile function conditional on known covariates and is often thought of as a generalization of estimating the conditional median (Koenker and Hallock, 2001). Just as conditional mean and conditional median estimation can be accomplished by using an appropriate loss function (sum of squares and absolute differences, respectively), estimating conditional quantiles

can be accomplished by minimizing the pinball loss defined as follows:

$$L_{\gamma}(z,q) := \begin{cases} (1-\gamma)(q-z), & z < q, \\ -\gamma(z-q), & z \ge q. \end{cases}$$
(37)

(Koenker, 2005; Steinwart and Christmann, 2011). As such, estimating the quantile surface can be framed as a risk-minimization problem, leaving only standard modeling choices to fill in for q in (37). Although initial efforts were focused on linear parametric quantile regressors (Koenker and Bassett Jr, 1978), in recent years, modeling efforts have focused on nonparametric varieties. Meinshausen (2006) adapted random forests to quantile regression. Takeuchi et al. (2006) leveraged Reproducing Kernel Hilbert Spaces to construct smooth quantile regressors. Closer to our application, Masserano et al. (2023) used neural networks to optimize the pinball loss to learn the quantile surface for their application.

Using the design points in the Berger–Boos set as sampled via the VGS or Polytope samplers, Algorithm 2 shows how we sample from the test statistic distribution defined at each design point to define a data set to fit a quantile regressor. We use the Gradient Boosting Regressor implemented in scikit-learn (Pedregosa et al., 2011) with the "quantile" loss function (i.e., pinball loss defined in eq. (37)) to fit the quantile surface for our numerical examples in Section 6. This algorithm involves a collection of hyperparameters (i.e., the minimum number of samples required to split an internal node, the minimum number of samples required to be a leaf node, the maximum depth of any individual estimator, the learning rate, and the number of estimators) which we determine using 10-fold cross validation in a pilot study ahead of our simulation experiments in Section 6. Although one may use any quantile regression approach to estimate the quantile surface, we emphasize the importance of choosing an approach that can accommodate a nonlinear surface in the parameters (such as the Gradient Boosting Regressor) as the parameter constraints are known to produce nonlinear quantile surfaces in even simple examples as seen in Batlle et al. (2023).

6 Numerical experiments

For scenarios within the linear-Gaussian case of data-generating process (19), the OSB interval can be regarded as the previous the state-of-the-art option for computing constraint-aware confidence intervals. Although these intervals have empirically achieved nominal coverage in applications (Patil et al., 2022; Stanley et al., 2022), they generally do not guarantee coverage (Batlle et al., 2023). As such, we use the OSB interval in the following numerical experiments as a main comparison point to the intervals defined in this paper. In scenarios where the OSB interval achieves at least nominal coverage, we show that our intervals are either competitive or better in terms of expected interval length. In scenarios where the OSB interval does not achieve nominal coverage, our intervals do achieve nominal coverage and can have shorter expected length. We provide four numerical experiments to make these points. The first set of two uses a constrained Gaussian noise model setup in two or three dimensions. These two experiments illustrate the aforementioned points in addition to constituent parts of the interval computation process due to the relatively low dimensions. The second set of two considers a wide-bin deconvolution setup inspired by particle unfolding in highenergy physics (Stanley et al., 2022). This setup features an 80-dimensional parameter space with a rank-deficient forward model and is thus a substantially more complicated computational scenario compared to the two or three-dimensional Gaussian noise models. These examples demonstrate the superior performance of our intervals over OSB in terms of both coverage and expected length.

In the following experiments, we form 68% confidence intervals, set $\eta = 0.01$ and compute γ

according to Lemma 3.1. We draw 10^3 observations from each data-generating process to estimate both coverage and expected interval length for our four interval constructions and the OSB interval. We additionally provide 95% confidence intervals in the form of orange line segments to characterize statistical error for both coverage and expected length estimates. The coverage confidence intervals are Clopper-Pearson intervals for the success probability parameter of a binomial distribution, while the expected length confidence intervals are the average length plus/minus the appropriately scaled standard error of the mean.

6.1 Constrained Gaussian in two dimensions

The two-dimensional Gaussian noise model is defined as follows:

$$\boldsymbol{y} = \boldsymbol{x}^* + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_2), \quad \boldsymbol{x}^* \in \mathbb{R}^2_+,$$
(38)

where $\varphi(\boldsymbol{x}) = x_1 - x_2$ and $\boldsymbol{x}^* = \begin{pmatrix} 0.5 & 0.5 \end{pmatrix}^\top$. The LLR is then given as follows:

$$\lambda(\mu, \boldsymbol{y}) = \min_{\substack{x_1 - x_2 = \mu \\ \boldsymbol{x} \in \mathbb{R}^2_+}} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2 - \min_{\boldsymbol{x} \in \mathbb{R}^2_+} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2.$$
(39)

This example first appeared in Tenorio et al. (2007) as a case where the OSB interval allegedly fails to achieve nominal coverage when the true parameter \boldsymbol{x}^* is such that $\varphi(\boldsymbol{x}^*) = 0$. However, Batlle et al. (2023) overturned this result by proving OSB validity in this case. As such, this example is important to include because of its historical context and OSB interval validity. The proof that the OSB interval covers in this particular example relies upon showing that $Q_{1-\alpha}^{\max} = \chi_{1,\alpha}^2$ for all $\alpha \in (0,1)$, where $\chi_{1,\alpha}^2$ is the upper α -quantile of a chi-squared distribution with one degree of freedom. Alternatively stated, it holds that $\lambda(\varphi(\boldsymbol{x}^*), \boldsymbol{y}; \mathbb{R}^2_+)$ is stochastically dominated by χ_1^2 . This result is shown in Lemma 4.4 of Batlle et al. (2023).

Estimated coverage and length results are shown in Figure 6.1. We note that all four of our interval constructions are competitive with OSB in terms of coverage, while all of our interval constructions have higher estimated expected length, apart from the Sliced constructions which are within statistical error of OSB. Since the OSB interval is defined using $Q_{1-\alpha}^{\max} = \chi_{1,\alpha}^2$ and the α -quantile surface of the LLR rapidly approaches this global max-quantile as one moves away from the origin (see Figure 5.3 in Batlle et al. (2023)), the OSB interval lengths are difficult to beat in practice with intervals based on the Berger–Boos sets since these sets likely contain parameter settings with quantiles near $\chi^2_{1,\alpha}$. The left panel of Figure 6.2 shows four realizations of the data-generating process with the observations shown as red points. For each observation, the blue points show uniformly distributed draws within its Berger–Boos set, sampled using the VGS sampler. Cross-referencing the spread of the Berger-Boos set samples in Figure 6.2 with Figure 5.3 in Batlle et al. (2023), it is clear that there are always samples in the parameter space where the quantile surface is nearly the same as the $\chi^2_{1,\alpha}$ quantile. Further, when including the Berger–Boos set in the interval construction, we instead construct our intervals using the γ -quantile, where $\gamma < \alpha$, as the LLR cutoff, resulting in a more relaxed constraint. This fact can be clearly observed in the central panel of Figure 6.2, showing the sampled γ -quantiles within the Berger-Boos set of one observation from the data generating process. Since a non-trivial portion of this distribution is above $\chi^2_{1,\alpha}$, the longer average length of the Global intervals is explained. In the right panel of Figure 6.2, for the same observation, we show the estimated sliced max-quantile function, $\widehat{m}_{\gamma}(\mu)$, in orange alongside $\chi^2_{1,\alpha}$. Since this estimated function is above $\chi^2_{1,\alpha}$ at their intersection points with the underlying LLR function shown in the solid blue line, it further makes sense that the Sliced interval constructions provide no additional length improvement compared to the OSB interval in this particular setting.



Figure 6.1: Estimated coverages and expected lengths across all four interval constructions and OSB for comparison at the 68% level for the two-dimensional constrained Gaussian setting. All four of our interval constructions are comparable to OSB with respect to coverage, but OSB shows better-expected length performance, aside from our two sliced interval constructions. Although the OSB intervals are defined using the global max-quantile ($Q_{1-\alpha}^{\max}$) and therefore can potentially be improved upon by limiting the considered parameter space via the Berger–Boos set, due to the rapidity with which the α -quantile surface meets the $\chi^2_{1,\alpha}$ quantile (see Figure 5.3 in Batlle et al. (2023)), the OSB interval lengths are difficult to beat in practice.

6.2 Constrained Gaussian in three dimensions

As seen in the previous example, in a case where the OSB interval is known to achieve nominal coverage, its expected length can be difficult to beat. However, OSB coverage guarantee can be difficult to prove or disprove, since it amounts to proving stochastic dominance on the non-trivial LLR statistic. In such situations, our intervals immediately provide a clear theoretical advantage. One such case involving a three-dimensional constrained Gaussian case was explored in Batlle et al. (2023), to which we now apply our four interval constructions. The three-dimensional Gaussian noise model is defined as follows:

$$\boldsymbol{y} = \boldsymbol{x}^* + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_3), \quad \boldsymbol{x}^* \in \mathbb{R}^3_+,$$
(40)

where $\varphi(\boldsymbol{x}) = x_1 + x_2 - x_3$ and $\boldsymbol{x}^* = \begin{pmatrix} 0.03 & 0.03 & 1 \end{pmatrix}^{\top}$. The LLR is then defined as follows:

$$\lambda(\mu, \boldsymbol{y}) = \min_{\substack{x_1 + x_2 - x_3 = \mu \\ \boldsymbol{x} \in \mathbb{R}^3_+}} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2 - \min_{\boldsymbol{x} \in \mathbb{R}^3_+} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2.$$
(41)

Figure 6.3 shows estimated coverage and expected length across all four interval constructions and the OSB interval. While the OSB interval fails to attain nominal coverage in this example, all four of our interval constructions do, with the Sliced constructions providing the best calibration. While the Global constructions pay a fairly steep price for coverage in expected interval length, the Sliced constructions navigate the trade-off well, paying for coverage with only slightly longer intervals compared to the OSB interval.

The setting of \boldsymbol{x}^* used here is slightly different than that of Batlle et al. (2023), where $\boldsymbol{x}^* = \begin{pmatrix} 0 & 0 & 1 \end{pmatrix}^\top$ was used. In Batlle et al. (2023), this setting was used as a counter-example for OSB coverage, since $Q_{\boldsymbol{x}^*}(1-\alpha) > \chi^2_{1,\alpha}$ for at least some $\alpha \in (0,1)$. However, as shown in Figure 5.5 of Batlle et al. (2023), when $\alpha = 0.05$, $Q_{\boldsymbol{x}^*}(1-\alpha) \le \chi^2_{1,\alpha}$ for $\boldsymbol{x}^* = \begin{pmatrix} t & t & 1 \end{pmatrix}^\top$ when t is approximately greater than $e^{-2} \approx 0.135$, which indicates that parameter settings violating stochastic dominance by χ^2_1 exist close to the parameter constraint boundary. The location of these key parameter settings



Figure 6.2: (Left) Four realizations of the data-generating process where the observations are shown in red. For each realization, the blue points are uniformly distributed samples from its Berger-Boos set, sampled using the VGS sampler. (Center) For a realization of the data-generating process, we plot the distribution of γ -quantiles for the points sampled by the VGS sampler. Notably, a non-trivial percent of these are above $\chi^2_{1,\alpha}$ defining the OSB interval. (Right) For the same realization, we plot the estimated sliced maxquantile function, $\hat{m}_{\gamma}(\mu)$ in orange alongside $\chi^2_{1,\alpha}$ in red. The blue points correspond to sampled parameter values, each of which has a functional and quantile value, while the solid blue line shows the LLR over the functional varies. All intervals can be read immediately from this image by inspecting where the blue LLR curve intersects the sampled points.



Figure 6.3: Estimated coverage and expected length across all four interval constructions and OSB for comparison at the 68% level for the three-dimensional constrained Gaussian example. All four of our interval constructions achieve nominal coverage while the OSB interval does not. While the Global interval constructions pay a steep price in expected length compared to OSB, the Sliced constructions are only slightly longer than OSB.

presents a challenge for the Polytope sampler described by Algorithm 4. In Appendix C.1, Algorithm 5 presents a modified version of Algorithm 4 that better handles sampling in this example.

Berger-Boos set experiment. For this model, we investigate the effect of changing the parameter η that controls the Berger-Boos construction in the global interval. We compute intervals for different \boldsymbol{y} and fixed $\boldsymbol{x}^* = (2, 2, 0), \boldsymbol{x}^* = (3, 3, 0)$ and $\boldsymbol{x}^* = (5, 5, 0)$ as η ranges between 0 and $\alpha = 0.32$. The lengths of such intervals, averaged over the data \boldsymbol{y} , are shown in Figure 6.4. In this example the maximum quantile is achieved at $\boldsymbol{x} = (0, 0, t)$ for large t, and, as expected, the benefit of using a small $\eta > 0$ becomes more pronounced as the true \boldsymbol{x}^* becomes farther from the point that achieves the maximum quantile. However, as η grows too close to α the downside of optimizing

the $1 - \alpha + \eta$ quantile instead of $1 - \alpha$ outweights the benefit of optimizing it over a smaller set, resulting in larger intervals. This suggests that a small $\eta > 0$ is a reasonable default, as suggested originally by Berger and Boos (1994).



Figure 6.4: Confidence interval lengths in the Berger–Boos setting, averaged over values of \boldsymbol{y} , for varying η and \boldsymbol{x}^* . The minimum average length occurs at a small $\eta > 0$, showing that the construction is beneficial if η is tuned correctly. This occurs because even for moderately small η , the Berger–Boos set, which is a three-dimensional sphere intersected with the non-negative orthant, avoids the point with the highest $1 - \alpha$ quantile.

6.3 Wide-bin deconvolution

While the numerical experiments in Sections (6.1) and (6.2) show that our interval constructions are competitive with the OSB interval in scenarios where it is known to provide coverage and superior to the OSB interval by achieving nominal coverage when OSB does not, this section shows the superior performance of our interval constructions relative to OSB in a more complex high-dimensional setting. We consider the problem of computing a confidence interval for the sum of adjacent bins of a deconvolved histogram as described in Stanley et al. (2022). This problem is a core statistical problem of particle unfolding in high-energy physics. For more detailed information, we refer the reader to Kuusela and Panaretos (2015); Kuusela (2016); Kuusela and Stark (2017); CMS Collaboration (2016, 2019).

The data-generating process is linear with Gaussian noise,

$$\boldsymbol{y} = \boldsymbol{K}\boldsymbol{x}^* + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}), \quad \boldsymbol{x}^* \ge \boldsymbol{0},$$
(42)

where $\mathbf{K} \in \mathbb{R}^{40 \times 80}$ and $\varphi(\mathbf{x}) = \mathbf{h}^{\top} \mathbf{x}$. The vector \mathbf{h} defines the bin-adjacent aggregation. In particle unfolding, the vectors \mathbf{x}^* and \mathbf{y} represent particle counts within discretized bins. A collection of \mathbf{h} vectors can act to sum the contents of adjacent bins to effectively lower the resolution of the inference problem. The LLR is then defined as follows:

$$\lambda(\mu, \boldsymbol{y}) = \min_{\substack{\boldsymbol{h}^{\top} \boldsymbol{x} = \mu\\ \boldsymbol{x} \in \mathbb{R}^{80}_{\perp}}} \|\boldsymbol{y} - \boldsymbol{K} \boldsymbol{x}\|_{2}^{2} - \min_{\boldsymbol{x} \in \mathbb{R}^{80}_{+}} \|\boldsymbol{y} - \boldsymbol{K} \boldsymbol{x}\|_{2}^{2}.$$
(43)

We emphasize two features of this setup that complicate the task of computing confidence intervals for $\varphi(\boldsymbol{x})$. First, the forward model, \boldsymbol{K} , has a non-trivial null space and a large condition number, making any inverse problem point estimation and UQ markedly challenging. Typically, this sort of ill-posedness is handled with regularization of some kind, but as is well-known in the inverse problem literature and specifically shown in Kuusela (2016), including such regularization induces a bias, which can undercut desired statistical guarantees (e.g., coverage) of the inference object of interest. Including constraints and focusing on a particular functional of the parameter vector *implicitly* regularizes the problem (Patil et al., 2022; Stanley et al., 2022), but shifts the problem difficulty to inference with constraints. Although Batlle et al. (2023) and this paper proposes a theoretical framework to perform inference with constraints, the second challenge is in the practical implementation due to the high-dimensional parameter space in scenarios like this example. As we show in the following sections, the Polytope sampler described by Algorithm 4 and quantile regression do an adequate job producing samples and fitting quantile surfaces in this high-dimensional space to ensure the desired coverage of the interval constructions.

As extensively discussed in Stanley et al. (2022), the OSB interval (i.e., using $\chi^2_{1,0.05}$ in the optimization-based interval construction) produces empirically valid confidence intervals in all tested scenarios, albeit typically with over-coverage. In the scenarios considered in Stanley et al. (2022), the underlying function generating the true histogram means ($\boldsymbol{x}^* \in \mathbb{R}^{80}_+$) was relatively smooth, likely contributing to the over-coverage. As such, we present two true parameter settings for \boldsymbol{x}^* in (42) to highlight two advantages of our interval constructions over the OSB interval. First, we use the original smooth parameter setting from Stanley et al. (2022) to show how our intervals improve over-coverage relative to the OSB interval by reducing the expected interval length. Second, we present an "adversarial" setting where our interval constructions achieve nominal coverage while the OSB interval does not. Figure 6.5 shows the smooth and adversarial settings for \boldsymbol{x}^* . We constructed the adversarial setting by first computing our interval constructions on the smooth setting and then looking at the maximum out-of-sample predicted quantile for a generated observation with a large predicted quantile. For each observation drawn within both settings, we draw 2.1×10^4 samples using the Polytope sampler as described by Algorithm 4.



Figure 6.5: Parameter values for the smooth and adversarial settings for x^* used to illustrate our interval construction versus the OSB interval. The adversarial setting is made more difficult by the sharp jumps in parameter values.

6.3.1 Smooth setting

Using the smooth x^* shown in Figure 6.5, Stanley et al. (2022) showed that the OSB interval over-covers at the 95% level. Furthermore, it was shown that the OSB interval was the shortest across a range of other interval options, including SSB, prior optimized, and minimax. As such, for this setting, we show that our interval constructions not only achieve nominal coverage, but the sliced constructions dramatically reduce over-coverage compared to OSB by producing substantially shorter intervals on average. Estimated coverage and expected interval lengths are shown in Figure 6.6. Both Global interval constructions and OSB dramatically over-cover which highlights the conservatism of the Global constructions. These estimated coverage values indicate that within each observation's Berger–Boos set, there is a parameter setting against which the method has to protect that is substantially more difficult to cover than the true realistic parameter setting. Interestingly, both Global constructions produce markedly longer intervals on average compared with the OSB interval. Aligning with the intuition from the Global and Sliced construction definitions, the Sliced intervals are less conservative as seen by their lower over-coverage and significantly smaller average lengths. Importantly, both Sliced constructions are shorter on average compared to the OSB interval, with the Sliced Inverted showing an 18.7% reduction in average length and the Sliced Optimized showing a 11.1% reduction in average length.



Figure 6.6: Estimated coverage and expected length across all four interval constructions and OSB at the 68% level for the **smooth** wide-bin deconvolution experiment. While the Global interval constructions overcover like the OSB interval, the Sliced interval constructions reduce both over-coverage and expected interval length.

6.3.2 Adversarial setting

The key result in Section 6.3.1 is that the Sliced interval constructions both reduce over-coverage and expected interval length compared with the OSB interval. In this section, we show that for the adversarial parameter setting the OSB interval does not achieve nominal coverage, whereas all four of our interval constructions do achieve nominal coverage while still reducing the expected interval length in the case of the Sliced interval constructions compared to the OSB interval. The corresponding estimated coverage and expected length results are shown in Figure 6.7.

Both Global interval constructions and the Sliced Optimized interval over-cover, with the Sliced Optimized over-covering to a lesser extent than the Global intervals. The Sliced Inverted interval achieves nominal coverage within the statistical uncertainty. The estimated expected lengths tell a story similar to that of the smooth example, with the Global intervals showing the longest average interval lengths, the Sliced intervals showing the shortest, and the OSB interval being between the two. Importantly, the Sliced intervals are again significantly shorter than the OSB interval, even though the OSB interval does not achieve nominal coverage. The Sliced Inverted interval shows a 18.9% average interval length reduction over OSB while the Sliced Optimized interval shows an 11.4% average interval length reduction.

For both the smooth and adversarial experiments, we note the differences in coverage and expected length between the two Sliced interval constructions. Although both approaches theoretically compute the same interval, significant differences can arise from their construction. For instance, the



Figure 6.7: Estimated coverage and expected length across all four interval constructions and OSB at the 68% level for the **adversarial** wide-bin deconvolution experiment. While the OSB interval fails to achieve nominal coverage, all four of our interval constructions do. Interestingly, the Sliced interval constructions are meaningfully shorter than the OSB interval while also providing coverage.

Sliced Inverted intervals are constructed by accepting individual functional values, which makes each point's acceptance dependent upon the quality of the quantile regressor at that point. By contrast, since the Sliced Optimized intervals essentially smooth over the max quantiles as a function of the functional space, the intervals are less sensitive to the quantile regressor's performance at any individual point. In the above simulation studies, there were realizations of the data for which the Sliced Inverted construction only accepted a single functional value sample thereby making the confidence interval a single point. We found the Sliced Optimized construction to be more robust in these settings as the max quantile values were shared in a sliding window over the functional space.

7 Conclusion

This paper proposes several confidence interval constructions for functionals in constrained illposed inverse problems. Our approach is based on two key ideas: data-adaptive constraints using a Berger–Boos construction and sampling-based inversion. Two independent decisions when constructing intervals provide four different valid intervals: Global versus Sliced, using the quantile function either over the entire or along level-set slices of the Berger-Boos set and Inverted versus Optimized, constructing the interval either by individually accepted functional values via the estimated quantile function or using the estimated quantile function in endpoint optimizations. All of the constructions are built upon the preliminary constraint by the data-informed Berger-Boos set, followed by a sampling procedure to estimate a quantile function that can be used to invert or optimize the interval endpoints. We have validated the method (including all four aforementioned interval constructions) through several numerical examples, demonstrating its ability to provide correct coverage, better calibration, and comparable or shorter interval length compared to the OSB interval baseline. Overall, our approach offers a flexible framework that can incorporate constraints directly and can be tailored to various types of inverse problems. The main takeaway is that data-adaptive constraining helps improve the length of the resulting confidence intervals, and enables sampling which makes it feasible to carry out the test inversion needed to construct confidence intervals with a desired nominal coverage.

There are several promising directions for future work. One direction is to find ways to extend our method to even higher-dimensional problems, which are more challenging. This would involve developing improved techniques to handle the curse of dimensionality and exploring the trade-off between accuracy and computational complexity. For the approaches in this paper in particular, this extension would require a more tailored sampling approach. Another direction is to leverage more sophisticated machine learning algorithms (deep learning models or ensemble methods) to improve the estimate of the quantile function and thus improve the accuracy and efficiency of our confidence intervals. Additionally, applying our approach to other applications involving ill-posed inverse problems, such as medical imaging or geophysics, would provide further validation of the effectiveness of our approach. Finally, it is of interest to conduct a further theoretical analysis of our approach under different constraints and noise conditions to better understand its limitations and strengths. This would involve studying the statistical properties of our confidence intervals and investigating the impact of various assumptions on their performance. Of particular interest are relaxations of the linear forward model and Gaussian noise assumptions to extend our method's application domain. Although the original theoretical foundation developed in Batlle et al. (2023) does not make these assumptions, our implementation relies upon them for the tractability of the sampling algorithms. Overall, these future research directions have the potential to demonstrate the applicability and robustness of our approach in a wide range of domains.

Acknowledgements

We thank members of the STAMPS (Statistical Methods for the Physical Sciences) Research Center at Carnegie Mellon University for fruitful discussions on this work. MS and MK were partially supported by NSF grants DMS-2053804 and PHY-2020295, JPL RSAs No. 1670375, 1689177 & 1704914 and a grant from the C3.AI Digital Transformation Institute. PB and HO acknowledge support from the Air Force Office of Scientific Research under MURI awards number FA9550-20-1-0358 (Machine Learning and Physics-Based Modeling and Simulation), FOA-AFRL-AFOSR-2023-0004 (Mathematics of Digital Twins) and by the Department of Energy under award number DE-SC0023163 (SEA-CROGS: Scalable, Efficient and Accelerated Causal Reasoning Operators, Graphs and Spikes for Earth and Embedded Systems). Additionally HO acknowledge support from the DoD Vannevar Bush Faculty Fellowship Program.

References

- Ashton, G. et al. (2022). Nested sampling for physical scientists. *Nature Reviews Methods Primers*, 2.
- Batlle, P., Stanley, M., Patil, P., Kuusela, M., and Owhadi, H. (2023). Optimization-based frequentist confidence intervals for functionals in constrained inverse problems: Resolving the burrus conjecture. arXiv preprint arXiv:2310.02461.
- Berger, R. L. and Boos, D. D. (1994). P values maximized over a confidence set for the nuisance parameter. *Journal of the American Statistical Association*, 89:1012–1016.
- Boyd, S. and Vandenberghe, L. (2004). Convex Optimization. Cambridge University Press.
- Buchner, J. (2023). Nested sampling methods. arXiv preprint arXiv:2101.09675v4.
- Burrus, W., Commission, U. A. E., and Laboratory, O. R. N. (1965). Utilization of a Priori Information by Means of Mathematical Programming in the Statistical Interpretation of Measured Distributions. Oak Ridge National Laboratory.

- Carpenter, J. (1999). Test inversion bootstrap confidence intervals. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 61(1):159–172.
- Casella, G. and Berger, R. L. (2002). Statistical Inference. Duxbery, 2 edition.
- Chen, Y., Dwivedi, R., Wainwright, M. J., and Yu, B. (2018). Fast mcmc sampling algorithms on polytopes. *Journal of Machine Learning Research*, 19:1–86.
- CMS Collaboration (2016). Measurement of differential cross sections for Higgs boson production in the diphoton decay channel in pp collisions at $\sqrt{s} = 8$ TeV. The European Physical Journal C, 76(13).
- CMS Collaboration (2019). Measurement of inclusive and differential Higgs boson production cross sections in the diphoton decay channel in proton-proton collisions at $\sqrt{s} = 13$ TeV. Journal of High Energy Physics, 2019(183).
- Dalmasso, N., Izbicki, R., and Lee, A. (2020). Confidence sets and hypothesis testing in a likelihoodfree inference setting. In *International Conference on Machine Learning*, pages 2323–2334. PMLR.
- Dalmasso, N., Masserano, L., Zhao, D., Izbicki, R., and Lee, A. B. (2024). Likelihood-free frequentist inference: Bridging classical statistics and machine learning for reliable simulator-based inference. *Electronic Journal of Statistics*, pages 5045–5090.
- Fisher, E., Schweiger, R., and Rosset, S. (2020). Efficient construction of test inversion confidence intervals using quantile regression. *Journal of Computational and Graphical Statistics*, 29(1):140–148.
- Garthwaite, P. H. and Buckland, S. T. (1992). Generating monte carlo confidence intervals by the robbins-monro process. *Journal of the Royal Statistical Society*, 41(1):159–171.
- Geng, X. and Xie, L. (2019). Data-driven decision making in power systems with probabilistic guarantees: Theory and applications of chance-constrained optimization. *Annual reviews in control*, 47:341–363.
- Gouriéroux, C., Holly, A., and Monfort, A. (1982). Likelihood ratio test, wald test, and kuhn-tucker test in linear models with inequality constraints on the regression parameters. *Econometrica*, 50:63–80.
- Höcker, A. and Kartvelishvili, V. (1996). Svd approach to data unfolding. Nuclear Instruments and Methods in Physics Research A, 372:469–481.
- Kannan, R. and Narayanan, H. (2012). Random walks on polytopes and an affine interior point method for linear programming. *Mathematics of Operations Research*, 37:1 20.
- Kibzun, A. I. and Kan, Y. S. (1997). Stochastic programming problems with probability and quantile functions. *Journal of the Operational Research Society*, 48(8):849–849.
- Koenker, R. (2005). Quantile regression, volume 38. Cambridge university press.
- Koenker, R. and Bassett Jr, G. (1978). Regression quantiles. Econometrica: journal of the Econometric Society, pages 33–50.

- Koenker, R. and Hallock, K. F. (2001). Quantile regression. *Journal of Economic Perspectives*, 15(4):143–156.
- Kuusela, M. (2016). Uncertainty quantification in unfolding elementary particle spectra at the Large Hadron Collider. PhD thesis, École Polytechnique Fédérale de Lausanne.
- Kuusela, M. and Panaretos, V. M. (2015). Statistical unfolding of elementary particle spectra: Empirical bayes estimation and bias-corrected uncertainty quantification. *The Annals of Applied Statistics*, 9(3):1671–1705.
- Kuusela, M. and Stark, P. B. (2017). Shape-constrained uncertainty quantification in unfolding steeply falling elementary particle spectra. *The Annals of Applied Statistics*, 11(3):1671–1710.
- Liu, J., Bowman, K., Meemong, L., et al. (2016). Carbon monitoring system flux estimation and attribution: impact of ACOS-GOSAT X_{CO_2} sampling on the inference of terrestrial biospheric sources and sinks. *Tellus B: Chemical and Physical Meteorology*, 66.
- Lovasz, L. (1999). Hit-and-run mixes fast. Mathematical Programming, 86:443 461.
- Lovasz, L. and Vempala, S. (2006). Hit-and-run from a corner. SIAM Journal on Computing, 35:985 1005.
- Masserano, L., Dorigo, T., Izbicki, R., Kuusela, M., and Lee, A. (2023). Simulation-based inference with waldo: Confidence regions by leveraging prediction algorithms or posterior estimators for inverse problems. *AISTATS*.
- Masserano, L., Shen, A., Doro, M., Dorigo, T., Izbiki, R., and Lee, A. (2024). Classification under nuisance parameters and generalized label shift in likelihood-free inference. *arXiv preprint* arXiv:2402.05330.
- Meinshausen, N. (2006). Quantile regression forests. Journal of Machine Learning Research, 7:983– 999.
- Molenberghs, G. and Verbeke, G. (2007). Likelihood ratio, score, and wald tests in a constrained parameter space. *The American Statistician*, 61(1):22–27.
- Montel, N. A., Alvey, J., and Weniger, C. (2023). Scalable inference with autoregressive neural ratio estimation. arXiv preprint arXiv:2308:08597v1.
- Narayanan, H. (2016). Ranomized interior point methods for sampling and optimization. The Annals of Applied Probability, 26:597–641.
- Ok, E. A. (2007). *Real Analysis with Economic Applications*. Econometric Society Monographs. Princeton University Press, Princeton, NJ, 1st edition. Includes bibliographical references and index.
- Pakman, A. and Paninski, L. (2014). Exact hamiltonian monte carlo for truncated multivariate gaussians. *Journal of Computational and Graphical Statistics*, 23:518 542.
- Panaretos, V. M. (2016). Statistics for Mathematicians. Springer International Publishing.
- Patil, P., Kuusela, M., and Hobbs, J. (2022). Objective frequentist uncertainty quantification for atmospheric CO₂ retrievals. SIAM/ASA Journal on Uncertainty Quantification, 10.

- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Pena-Ordieres, A., Luedtke, J. R., and Wachter, A. (2020). Solving chance-constrained problems via a smooth sample-based nonlinear approximation. SIAM Journal on Optimization, 30(3):2221– 2250.
- Robertson, T., Wright, F. T., and Dykstra, R. (1988). Order Restricted Statistical Inference. Wiley.
- Rodgers, C. D. (2000). Inverse Methods for Atmospheric Sounding. World Scientific Publishing.
- Rust, B. W. and Burrus, W. R. (1972). Mathematical Programming and the Numerical Solution of Linear Equations. American Elsevier.
- Rust, B. W. and O'Leary, D. P. (1994). Confidence intervals for discrete approximations to ill-posed problems. *Journal of Computational and Graphical Statistics*, 3(1):67–96.
- Schmitt, S. (2012). Tunfold, an algorithm for correcting migration effects in high energy physics. Journal of Instrumentation, 7:T10003.
- Shapiro, A. (1988). Towards a unified theory of inequality constrained testing in multivariate analysis. *International Statistical Review/Revue Internationale de Statistique*, pages 49–62.
- Skilling, J. (2004). Nested sampling. AIP Conference Proceedings, 735:295.
- Smith, R. L. (1984). Efficient monte carlo procedures for generating points uniformly distributed over bounded regions. Operations Research, 32:1296–1308.
- Stanley, M., Patil, P., and Kuusela, M. (2022). Uncertainty quantification for wide-bin unfolding: one-at-a-time strict bounds and prior-optimized confidence intervals. *Journal of Instrumentation*, 17.
- Stark, P. B. (1992a). Inference in infinite-dimensional inverse problems: discretization and duality. Journal of Geophysical Research: Solid Earth, 97(B10):14055–14082.
- Stark, P. B. (1992b). Inference in infinite-dimensional inverse problems: Discretization and duality. Journal of Geophysical Research, 97(B10):14055–14082.
- Steinwart, I. and Christmann, A. (2011). Estimating conditional quantiles with the help of the pinball loss. *Bernoulli*, 117(1):211–225.
- Takeuchi, I., Le, Q. V., Sears, T. D., and Smola, A. J. (2006). Nonparametric quantile estimation. Journal of Machine Learning, 7:1231–1264.
- Tenorio, L., Fleck, A., and Moses, K. (2007). Confidence intervals for linear discrete inverse problems with a non-negativity constraint. *Inverse problems*, 23(2):669.
- Voelker, A. R., Gosmann, J., and Stewart, T. C. (2017). Simultaneous confidence intervals for linear estimates of linear functionals. Technical report, Centre for Theoretical Neuroscience.
- Wolak, F. A. (1987). An exact test for multiple inequality and equality constraints in the linear regression model. *Journal of the American Statistical Association*, 82:782–793.

Wolak, F. A. (1989). Testing inequality constraints in linear econometric models. Journal of Econometrics, 41:205–235.

Appendix

This document serves as an appendix to the paper "Confidence intervals for functionals in constrained inverse problems via data-adaptive sampling-based calibration". Below we provide an organization for the appendix, followed by a summary of the main notation used in both the paper and the appendix. The equation and figure numbers in this appendix begin with the letter "S" to differentiate them from those appearing in the main paper.

Organization

The content of this appendix is organized as follows.

Appendix	Description
Appendix A	Proofs in Section 3 (Proofs of Lemma 3.1 and Corollary 3.2)
Appendix B	Proofs in Section 4 (Proof of Theorem 4.1)
Appendix C	Additional details and illustrations in Section 6 (importance-like sampler)

Table 2: Roadmap of the supplement.

Notation

An overview of the main notation used in this paper is as follows.

Notation	Description
Non-bold lower or upper case	Denotes scalars (e.g., α , μ , Q).
Bold lower case	Denotes vectors (e.g., $\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{h}$).
Bold upper case	Denotes matrices (e.g., K , I).
Calligraphic font	Denotes sets (e.g., $\mathcal{X}, \mathcal{C}, \mathcal{D}$).
\mathbb{R}	Set of real numbers.
\mathbb{R}_+	Set of non-negative real numbers.
[n]	Set $\{1, \ldots, n\}$ for a positive integer n .
$\ \boldsymbol{u}\ _2$	The ℓ_2 norm of vector \boldsymbol{u} .
$\ f\ _{L_2}$	The L_2 norm of function f .
$K^ op$	The transpose of a matrix $\boldsymbol{K} \in \mathbb{R}^{m \times p}$.
I_m or I	The $m \times m$ identity matrix.
$oldsymbol{v} \leq oldsymbol{u}$	Lexicographic ordering for vectors \boldsymbol{v} and \boldsymbol{u} .
$A \preceq B$	The Loewner ordering for symmetric matrices \boldsymbol{A} and \boldsymbol{B} .
$1{A}$	Indicator random variable associated with event A.
$Y = \mathcal{O}_{\alpha}(X)$	Deterministic big-O notation, indicating that Y is bounded by $ Y \leq C_{\alpha} X$.
C_{lpha}	A numerical constant that may depend on the parameter α in context.
${\mathcal O}_p$	Probabilistic big-O notation.
$\xrightarrow{\mathrm{p}}$	Convergence in probability.
$X \succeq Y$	Stochastic dominance of X by Y, indicating that $\mathbb{P}(X \ge z) \ge \mathbb{P}(Y \ge z)$ for all $z \in \mathbb{R}$.

Table 3: Summary of main notation used in the paper and the appendix.

A Proofs in Section 3

A.1 Proof of Lemma 3.1

We reproduce the original argument in Berger and Boos (1994), originally stated in terms of p-values, translated into our quantile setting. Fix any $x^* \in \mathcal{X}$ and consider the sets:

- $A_1 = \{ \boldsymbol{y} : B_{\eta}(\boldsymbol{y}) \ni \boldsymbol{x}^* \}$
- $A_2 = \{ \boldsymbol{y} : \lambda(\mu^*, \boldsymbol{y}) \le Q_{\boldsymbol{x}}(1-\gamma) \}$
- $A_3 = \{ \boldsymbol{y} : \lambda(\mu^*, \boldsymbol{y}) \leq \bar{q}_{\gamma, \eta}(\mu^*) \} = \{ \boldsymbol{y} : \lambda(\mu^*, \boldsymbol{y}) \leq \sup_{\boldsymbol{x} \in \mathcal{B}_{\eta} \cap \Phi_{\mu^*}} Q_{\boldsymbol{x}}(1-\gamma) \}$

We now have that $\mathbb{P}(\boldsymbol{y} \in A_1) = 1 - \eta$, $\mathbb{P}(\boldsymbol{y} \in A_2) = 1 - \gamma$, and that $(A_1 \cap A_2) \subset (A_1 \cap A_3)$. Therefore,

$$\mathbb{P}(\boldsymbol{y} \notin A_3) = \mathbb{P}(\boldsymbol{y} \notin A_3, \boldsymbol{y} \in A_1) + \mathbb{P}(\boldsymbol{y} \notin A_3, \boldsymbol{y} \notin A_1)$$
(S.1)

 $\leq \mathbb{P}(\boldsymbol{y} \notin A_2, \boldsymbol{y} \in A_1) + \mathbb{P}(\boldsymbol{y} \notin A_1)$ (S.2)

$$\leq \mathbb{P}(\boldsymbol{y} \notin A_2) + \mathbb{P}(\boldsymbol{y} \notin A_1) \tag{S.3}$$

$$\gamma + \eta$$
 (S.4)

so that $\mathbb{P}(\boldsymbol{y} \in A_3) \geq 1 - \gamma - \eta$. Imposing $1 - \gamma - \eta \geq 1 - \alpha$ gives the desired result.

A.2 Proof of Corollary 3.2

By definition, $\bar{q}^{\mu}_{\gamma,\eta} \leq \bar{q}_{\gamma,\eta}$ for all $\mu \in \mathbb{R}$. Therefore, $C^{\rm sl}_{\alpha}(\boldsymbol{y}; \mathcal{B}_{\eta}) \subseteq C^{\rm gl}_{\alpha}(\boldsymbol{y}; \mathcal{B}_{\eta})$, and thus

$$\mathbb{P}_{\boldsymbol{x}^*}\left(\mu^* \in C^{\mathrm{gl}}_{\alpha}(\boldsymbol{y}; \mathcal{B}_{\eta})\right) \ge \mathbb{P}_{\boldsymbol{x}^*}\left(\mu^* \in C^{\mathrm{sl}}_{\alpha}(\boldsymbol{y}; \mathcal{B}_{\eta})\right) \ge 1 - \alpha.$$
(S.5)

B Proof of Theorem 4.1

Throughout the proof, we make use of the following lemma:

Lemma B.1. Let $f : \mathbb{R}^n \to \mathbb{R}$ and $\mathcal{X} \subset \mathbb{R}^n$ such that $\min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ is achieved, and at least one of the minimizers \mathbf{x}^* satisfies:

- 1. f is continuous at x^* ,
- 2. \mathbf{x}^* is not an isolated point of \mathcal{X} (i.e., $\forall \delta > 0, B_{\delta}(\mathbf{x}^*) \cap \mathcal{X} \neq \emptyset$).

Let μ be a measure on \mathcal{X} such that $\mu(B) > 0$ for all $B \subseteq \mathcal{X}$ such that $\lambda_{Leb}(B) > 0$ (where $\lambda_{Leb}(B)$) refers here to the Lebesgue measure of the set B). Let $Y_m = \min_{i=1,...,m} f(\mathbf{x}_i)$, where \mathbf{x}_i are *i.i.d.* samples

from μ . Then $Y_m \xrightarrow{p} f(\boldsymbol{x}^*)$.

Proof. Fix $\varepsilon > 0$ and let us show that $\mathbb{P}(|Y_m - f(\boldsymbol{x}^*)| > \varepsilon) \to 0$. Since f is continuous at \boldsymbol{x}^* there exists a $\delta > 0$ such that $f(B_{\delta}(\boldsymbol{x}^*)) \subset B_{\varepsilon}(f(\boldsymbol{x}^*))$ so that,

$$\mathbb{P}(|Y_m - f(\boldsymbol{x}^*)| \ge \varepsilon) \le \mathbb{P}(\boldsymbol{x}_i \notin B_{\delta}(\boldsymbol{x}^*), \ \forall i = 1, \dots, m) = (\mathbb{P}_{\boldsymbol{x} \sim \mu}(\boldsymbol{x} \notin B_{\delta}(\boldsymbol{x}^*)))^m.$$
(S.6)

Since \boldsymbol{x}^* is not an isolated point, we have $\lambda_{\text{Leb}}(B_{\delta}(\boldsymbol{x}^*) \cap \mathcal{X}) > 0$ and therefore $\mathbb{P}_{\boldsymbol{x} \sim \mu}(\boldsymbol{x} \notin B_{\delta}(\boldsymbol{x}^*)) < 1$ and $\mathbb{P}(|Y_m - f(\boldsymbol{x}^*)| \geq \varepsilon) \to 0$.

We begin by proving that the empirical maximums of the quantiles obtained both by Algorithm 1 and Algorithm 2 (assuming the quantile regressor is consistent) converge to the true max quantile.

Algorithm 1 Let $\bar{q}_{\gamma,\eta}^{\text{de}} := \max_{i=1,\dots,M} \hat{q}_{\gamma}^{i}(N)$, where we explicitly write the dependence with the number of samples and the index *i* refers to the quantile estimated at the *i*-th sampled point \boldsymbol{x}_{i} . We aim to show that $\bar{q}_{\gamma,\eta}^{\text{de}} \xrightarrow{\mathbf{P}} \bar{q}_{\gamma,\eta}$. We know that $\hat{q}_{\gamma}^{i}(N)$ converges in probability to $Q_{P_{\boldsymbol{x}_{i}}}(1-\gamma)$ as $N \to \infty$. We have

$$\left| \bar{q}_{\gamma,\eta}^{de} - \bar{q}_{\gamma,\eta} \right| \le \left| \bar{q}_{\gamma,\eta}^{de} - \max_{i=1,\dots,M} Q_{P_{\boldsymbol{x}_{i}}}(1-\gamma) \right| + \left| \max_{i=1,\dots,M} Q_{P_{\boldsymbol{x}_{i}}}(1-\gamma) - \bar{q}_{\gamma,\eta} \right|.$$
(S.7)

The first term can be made smaller than $\varepsilon/2$ as $N \to \infty$ by convergence of the estimator, and the second term can be made smaller than $\varepsilon/2$ as $M \to \infty$ by application of Lemma B.1 to the quantile function (maximizing instead of minimizing).

Algorithm 2 The proof is identical to that of Algorithm 1, with the only difference of replacing the quantiles estimated via Monte Carlo sampling to those estimated by the quantile regression, and N to $M_{\rm tr}$, the number of samples needed to train the quantile regression. Since the quantile regression is assumed to be consistent, the first term can be made arbitrarily small as M_{tr} grows, and the result follows.

Since identical convergence results apply for both algorithms, henceforth we will not explicitly distinguish: The proof is written in terms of N, which can be replaced by M_{tr} .

B.1 Proof of Statement 1 (Global Inverted)

Recall,

$$C_{\text{inv}}^{\text{gl}}(\boldsymbol{y}) = \left[\min_{k \in \{1,\dots,M\}: \lambda(\varphi(\boldsymbol{x}_k), \boldsymbol{y}) \le \widehat{q}(N)} \varphi(\boldsymbol{x}_k), \max_{k: \lambda(\varphi(\boldsymbol{x}_k), \boldsymbol{y}) \le \widehat{q}(N)} \varphi(\boldsymbol{x}_k)\right]$$
(S.8)

$$= \left[\min_{k \in \{1, \dots, M\} : \lambda(\mu_k, \boldsymbol{y}) \le \widehat{q}(N)} \mu_k, \max_{k : \lambda(\mu_k, \boldsymbol{y}) \le \widehat{q}(N)} \mu_k \right]$$
(S.9)

where $\widehat{q}(N) := \max_{i=1,...,M} \widehat{q}_{\gamma}^{i}(N)$, the estimated quantiles of the sampled $\boldsymbol{x}_{i} \in \mathcal{B}_{\eta}$ and $\mu_{i} := \varphi(\boldsymbol{x}_{i})$. Note that μ_{i} are samples in $\varphi(\mathcal{B}_{\eta}) \subset \mathbb{R}$. Also note that we have more explicitly written out the interval definition (i.e., Equation (26)) to emphasize clarity rather than presentation.

Consider the left extreme of the interval, a similar argument follows from the right extreme. Consider three quantities:

$$\mu_1(N, M) := \min \mu_k \quad \text{s.t.} \quad k = 1, \dots, M \text{ and } \lambda(\mu_k, \boldsymbol{y}) \le \widehat{q}(N) \tag{S.10}$$

$$\mu_2(M) := \min \mu_k \quad \text{s.t.} \quad k = 1, \dots, M \text{ and } \lambda(\mu_k, \boldsymbol{y}) \le \bar{q}_{\gamma, \eta}$$
 (S.11)

$$\mu_3 := \min \mu \quad \text{s.t.} \quad \mu \in \varphi\left(\mathcal{B}_{\eta}\right) \text{ and } \lambda(\mu, \boldsymbol{y}) \leq \bar{q}_{\gamma,\eta}$$
(S.12)

Our goal is to show that as $N, M \to \infty$, $\mu_1 \xrightarrow{p} \mu_3$. Lemma B.1 shows that $\mu_2 \xrightarrow{p} \mu_3$. Indeed, the minimization over the indices k such that the condition is satisfied can be seen as a rejection sampling strategy in which all accepted samples are samples of the feasible region of the optimization in μ_3 . As M grows, since the sampler eventually samples all areas of $\varphi(\mathcal{B}_{\eta})$, some samples are guaranteed to be close to the optimum with high probability. Finally, for fixed M and N going to infinity, $\mu_1 \xrightarrow{p} \mu_2$. This follows from the continuity of the optimization problem with respect to the right-hand side of the constraint, and the fact that $\max_{i=1,\dots,M} q(\boldsymbol{x}_i) \xrightarrow{p} \bar{q}_{\gamma,\eta}$. It follows that as $N, M \to \infty, \mu_1 \xrightarrow{p} \mu_3$.

B.2 Proof of Statement 2 (Sliced Inverted)

The proof technique is similar to the one of Statement 1, replacing $\widehat{q}(N) := \max_{i=1,\dots,M} \widehat{q}^i_{\gamma}(N)$ for $\widehat{q}^k_{\gamma}(N)$. Defined then, similarly as in the previous proof:

 $\mu_1(N, M) := \min \mu_k \quad \text{s.t.} \quad k = 1, \dots, M \text{ and } \lambda(\mu_k, \boldsymbol{y}) \le \hat{q}^k_{\gamma}(N) \tag{S.13}$

$$\mu_2(M) := \min \mu_k \quad \text{s.t.} \quad k = 1, \dots, M \text{ and } \lambda(\mu_k, \boldsymbol{y}) \le \bar{q}_{\gamma, \eta}(\mu_k) \tag{S.14}$$

$$\mu_3 := \min \mu \quad \text{s.t.} \quad \mu \in \varphi(\mathcal{B}_\eta) \text{ and } \lambda(\mu, \boldsymbol{y}) \le \bar{q}_{\gamma, \eta}(\mu)$$
 (S.15)

where $\bar{q}_{\gamma,\eta}(\mu) = \max_{\boldsymbol{x}\in\Phi_{\mu}\cap\mathcal{B}_{\eta}} Q_{\boldsymbol{x}}(1-\gamma)$. As $N \to \infty$, $\hat{q}_{\gamma}^{k}(N) \to Q_{\boldsymbol{x}_{k}}(1-\gamma)$. Lemma B.1 can be used to show $\mu_{2} \xrightarrow{\mathrm{p}} \mu_{3}$, because the sampler strategy used that first samples \boldsymbol{x}_{k} and then accepts $\mu_{k} = \varphi(\boldsymbol{x}_{k})$ as a sample if $\lambda(\varphi(\boldsymbol{x}_{k}), \boldsymbol{y}) < q_{\gamma}^{k}$, it can be shown that for every feasible point μ , there is eventually a sample close to it.

Finally, μ_1 will become arbitrarily close to μ_2 as M grows large, since both are taking the minimum over samples that are sampled densely from the feasible set, meaning accepted μ_k will eventually be close to μ_3 for both the case of μ_1 and the case of μ_2

B.3 Proof of Statement 3 (Global Optimized)

We prove the continuity of the optimization problem,

$$\max_{\mu \in \varphi(\mathcal{B}_{\eta})} \mu \text{ s.t. } \lambda(\mu, \boldsymbol{y}) \le q,$$
(S.16)

as a function of q in the positive measure interval $(\lambda(\bar{\mu}, \boldsymbol{y}), \bar{q}_{\gamma,\eta}]$. Therefore, convergence in probability follows as we have convergence in probability to $\bar{q}_{\gamma,\eta}$ as $N, M \to \infty$, which in particular implies that the maximum quantile estimate is in the interval $(\lambda(\bar{\mu}, \boldsymbol{y}), \bar{q}_{\gamma,\eta}]$ almost surely. We do so by appealing to the maximum theorem (Ok, 2007, § E.3), which, in general, guarantees continuity of functions of the form $f^*(\theta) = \sup\{f(x,\theta) : x \in C(\theta)\}$ as long as f is continuous, C is a continuous compact-valued correspondence and $C(\theta)$ is non-empty for all $\theta \in \Theta$. The continuity of C comes from the continuity of the LLR, f is equal to the identity and the strict feasibility condition ensures C is non-empty in $\Theta := (\lambda(\bar{\mu}, \boldsymbol{y}), \bar{q}_{\gamma,\eta}]$.

B.4 Proof of Statement 3 (Sliced Optimized)

We will prove sufficient conditions for convergence of $\inf_{\mu:\widehat{f}_k(\mu)\geq 0} \mu$ to $\inf_{\mu:f(\mu)\geq 0} \mu$ as \widehat{f}_k converges to f, and the result will follow by taking $\widehat{f}_k(\mu) = \widehat{m}_{\gamma}(\mu) - \lambda(\mu, \boldsymbol{y})$ and $f(\mu) = m_{\gamma}(\mu) - \lambda(\mu, \boldsymbol{y})$. A similar argument can be repeated for the supremum. Use use the notation \widehat{f}_k to indicate that k sampled points are used to estimate this function via the definition of $\widehat{m}_{\gamma}(\mu)$ (see Section 3.2).

Lemma B.2. Let $f : \mathbb{R} \to \mathbb{R}$ be a function, and let f_k be a sequence of functions $f_k : \mathbb{R} \to \mathbb{R}$. Let $\mu^* = \inf_{f(\mu) \ge 0} \mu$ and $\mu_k = \inf_{f_k(\mu) \ge 0} \mu$. Let the sequence of functions $\{f_k\}$ be such that for all $\delta > 0$,

$$\mathbb{P}\left(\sup_{\mu}|f_k(\mu) - f(\mu)| > \delta\right) \to 0 \text{ as } k \to \infty,$$
(S.17)

namely, f_k converges in probability uniformly to f. Furthermore, let f be such that for all $\varepsilon > 0$, there exists $\delta' > 0$ such that $|f(\mu)| > \delta'$ if and only if $|\mu - \mu^*| < \varepsilon$. Then, we have $\mu_k \xrightarrow{p} \mu^*$.

We then have $\mu_k \xrightarrow{p} \mu^*$.

Proof. Assume for the sake of contradiction that there exists $\varepsilon > 0$ such that $\mathbb{P}(|\mu_k - \mu^*| \ge \varepsilon)$ does not go to 0. Then, by the condition on f, it must follow that $\mathbb{P}(|f(\mu_k)| > \delta)$ does not go to 0. But,

$$\mathbb{P}\left(|f(\mu_k)| > \delta\right) \le \mathbb{P}\left(|f(\mu_k) - f_k(\mu_k)| > \delta\right) + \mathbb{P}\left(|f_k(\mu_k)| > \delta\right),\tag{S.18}$$

and the right-hand side goes to 0 by uniform convergence in probability (first term) and by feasibility of μ_k (second term).

C Additional details and illustrations in Section 6

C.1 Importance-like sampler for the three-dimensional example in Section 6.2

Since the parameter settings with quantiles meaningfully larger than $\chi^2_{1,\alpha}$ are located close to the constraint boundary, a sampling challenge is presented. Using the samplers as described in Section 5.1 results in under-sampling of this large-quantile region since both samplers provide uniform random samples over the Berger-Boos set. Algorithm 5 presents a modified version of Algorithm 1, an importance-like sampler to increase the probability mass of samples close to the constraint boundary. Note, we say "importance-like" because we do not provide any theoretical guarantee regarding this sampler's ability to produce draws from a particular target distribution. We tailored Algorithm 5 to settings with a non-negativity constraint and hand-tuned the length scale parameter to the particular three-dimensional example in Section 6.2.

Algorithm 5 Importance-like sampler for three-dimensional constrained Gaussian

Input: Number of samples: $M \in \mathbb{N}$, inverse length scale: γ_p , and order of norm: $q \in (0, 1)$.

1: Instantiate a list S of length M to store sampled points.

2: while $|\mathcal{S}| < M$ do Draw $M - |\mathcal{S}|$ realizations from Algorithm 4: $\tilde{x}_1, \ldots, \tilde{x}_{M-|\mathcal{S}|}$ 3: 4: for k = 1, ..., M - |S|: do Compute the probability of accepting the k-th draw: $p_k := \exp\left(-\gamma_p \|\widetilde{\boldsymbol{x}}_k\|_q\right)$. 5:Draw $z_k \sim \text{Bernoulli}(p_k)$. 6: if $z_k = 1$ then 7: $\mathcal{S}[k] \leftarrow \widetilde{x}_k$ 8: 9: end if 10: end for 11: end while **Output**: Sampled parameters in Berger–Boos set \mathcal{S} .

The key to Algorithm 5 is the additional accept/reject step where the k-th sample is accepted with probability p_k . Additionally, by setting $q \in (0, 1)$, we prioritize retaining samples closer to the non-negativity boundary. This effect can be seen in Figure S.1, where the vast majority of samples are found closer to the constraint boundary. The ability of Algorithm 5 to better sample the high-quantile regions of the Berger-Boos set can be seen in Figure S.2. In particular, for the functional values $\mu \in (-4, -2)$, the importance-like sampler is substantially more effective than the Polytope sampler at finding parameter setting with γ -quantiles greater than 1.1.



Figure S.1: When sampling points using Algorithm 5, the vast majority of samples are found closer to the nonnegativity constraint boundary. The true parameter setting is shown by the red point, while the parameter settings sampled by Algorithm 5 are shown by the blue points. This sampling prioritization helps adequately sample the regions of the Berger–Boos set where the quantile surface is larger than $\chi^2_{1,\alpha}$. Furthermore, the vast majority of sampled points lie within the Berger–Boos set with some lying outside within the bounding polytope.



Figure S.2: The importance-like sampler described by Algorithm 5 is more effective than the Polytope sampler described by Algorithm 4 at sampling parameter settings with γ -quantile greater than 1.1. Each parameter setting sampled by Algorithm 5 is shown by a blue point. This improved ability helps ensure the coverage guarantee shown in the left panel of Figure 6.3.