

Carnegie Mellon University

Statistics & Data Science

Adaptively Calibrated Optimization- Based Confidence Intervals for Inverse Problem Uncertainty Quantification

Mike Stanley (CMU) | February 28, 2024 | SIAM UQ 2024 | Trieste, Italy

Collaborators



Pau Batlle (Caltech)



Pratik Patil (UC Berkeley)



Mikael Kuusela (CMU)




Houman Owhadi (Caltech)

Carnegie Mellon University
Statistics & Data Science

Caltech

Berkeley
UNIVERSITY OF CALIFORNIA

Setting of interest



- Deterministic forward model: $f : \mathbb{R}^p \rightarrow \mathbb{R}^n, \mathbf{x} \mapsto f(\mathbf{x})$, (e.g., $f(\mathbf{x}) = \mathbf{K}\mathbf{x}$, $\mathbf{K} \in \mathbb{R}^{n \times p}$)
- Forward model parameter constraints: $\mathbf{A}\mathbf{x} \leq \mathbf{b}$ (e.g., $\mathbf{x} \geq \mathbf{0}$)
- Additive noise: $\mathbf{y} = f(\mathbf{x}) + \boldsymbol{\varepsilon}$, $\boldsymbol{\varepsilon} \sim N(\mathbf{0}, \boldsymbol{\Sigma})$, (more generally, $\mathbf{y} \sim P_{\mathbf{x}}$)
- Inferential object(s): parameter functionals, $\varphi(\mathbf{x}) \in \mathbb{R}$ (e.g., $\varphi(\mathbf{x}) = \mathbf{h}^T \mathbf{x}$)
- Applications where this setting arises: carbon flux inversion [Stanley et al., 2024b], remote sensing (XCO₂) [Patil et al., 2022], and particle unfolding [Kuusela, 2016], [Stanley et al., 2022], [Han et al., 2023].

UQ in this setting and some challenges

Inverse Problem Uncertainty Quantification

=

Reporting **statistically guaranteed** uncertainty quantification of the inferred functional value following from the noisy observation and the forward model

Statistically guaranteed: a confidence interval, $I(\mathbf{y})$, with a **coverage** guarantee, i.e.,
 $\forall \mathbf{x}^* \in \mathcal{X}, \mathbb{P}(\varphi(\mathbf{x}^*) \in I(\mathbf{y})) \geq 1 - \alpha$ for a chosen level $\alpha \in [0,1]$.

- Ill-posed problems make $f^{-1}(\mathcal{E}(\mathbf{y}))$ difficult to work with (e.g., $\text{null}(\mathbf{K}) \neq \{\mathbf{0}\}$),
- Making $I(\mathbf{y})$ constraint-aware (e.g., $\mathbf{x} \geq \mathbf{0}$) while retaining the desired coverage guarantee is highly non-trivial.

Optimization-based confidence intervals provide a start to a solution

There is a way to frame the interval computation as two endpoint optimizations

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

such that

$$\forall \mathbf{x}^* \in \mathcal{X}, \mathbb{P} \left(\varphi(\mathbf{x}^*) \in I(\psi_\alpha^2, \mathbf{y}) \right) \geq 1 - \alpha$$

where

$$D(\psi_\alpha^2, \mathbf{y}) := \{ \mathbf{x} : \|\mathbf{y} - \mathbf{K}\mathbf{x}\|_2^2 \leq \psi_\alpha^2 \text{ and } \mathbf{A}\mathbf{x} \leq \mathbf{b} \}.$$

A key challenge:
setting ψ_α^2 to obtain
this coverage
guarantee

Optimization-based confidence intervals provide a start to a solution (cont.)

- They provide a **start to a solution** because they,
 - reframe inference as optimization (good for computation),
 - elegantly handle the parameter constraints in the endpoint optimizations.
- However, **setting ψ_α^2** to provide the coverage guarantee turns out to be non-trivial.
 - For **simultaneous (SSB) coverage** : $\psi_\alpha^2 := \chi_{n,\alpha}^2$ [Stark, 1992]
 - For **one-at-a-time (OSB) coverage** : $\psi_\alpha^2 := \chi_{1,\alpha}^2 + s^2$, where $s^2 = \min_{\mathbf{x}: \mathbf{Ax} \leq \mathbf{b}} \|\mathbf{y} - \mathbf{Kx}\|_2^2$
[Patil et al., 2022], [Rust and O’Leary, 1994], [Stanley et al., 2022]
- However, the OSB setting does not hold in general [Tenorio et al. 2007, Batlle et al. 2023]

An **outline** of this talk and some main **take-aways**

1. Building on the work of [Batlle et al. 2023], we present a method to set ψ_α^2 in a data-dependent way to achieve interval coverage and improve interval length relative to OSB.
 - **Take-away**: our method is the first computationally feasible approach to properly calibrate these optimization-based intervals.
 - We call it **adaOSB** for “adaptive OSB”
2. We explore three numerical studies to demonstrate the method and its advantages.
 - **Take-away**: our method provides coverage in low dimensional ($p = 3$) example where OSB does not, and improves interval length in a scenario where OSB empirically over-covers ($p = 80$).

The optimized interval can be seen an inverted hypothesis test

- There is a particular hypothesis test and log-likelihood ratio test statistic recovering the interval

$$\overset{\text{OSB interval}}{I(\psi_\alpha^2, \mathbf{y})} = \overset{\text{Inverted hypothesis test}}{\{\mu \in \mathbb{R} : \lambda(\mu, \mathbf{y}) \leq q_\alpha\}} \text{ where } \psi_\alpha^2 := q_\alpha + s^2$$

$\lambda(\mu, \mathbf{y})$ is an LLR test statistic and q_α is a critical value ensuring $\mathbb{P}(\text{type-1 error}) \leq \alpha$ for all $\mathbf{x}^* \in \mathcal{X}$

- **Theorem 2.4** [Batlle et al. 2023]: the critical value controlling type-1 error of the test can also be used to calibrate $I(\psi_\alpha^2, \mathbf{y})$

The previously mentioned q_α is both difficult to obtain and statistically conservative

- **Difficult to obtain:** finding q_α involves solving a chance constrained optimization which is known to be strong NP-hard and non-convex [Batlle et al. 2023].
- **Statistically conservative:** In order for q_α to control type-1 error for all $\mathbf{x}^* \in \mathcal{X}$, we are protecting against all potential true parameter states and therefore might be overly conservative.

We will demonstrate a method that **avoids solving the complicated optimizations** and provides **length benefits** in some cases

The hypothesis test connection can calibrate these optimization-based intervals

- First, let $Q_x : [0,1] \rightarrow \mathbb{R}$ be the quantile function of $\lambda(\mu, \mathbf{y})$ at \mathbf{x} , i.e. $\mathbb{P}(\lambda(\mu, \mathbf{y}) \leq Q_x(1 - \alpha)) = 1 - \alpha$.
- Let $\mathbf{x}^* \in \mathcal{X}$ denote the true but unknown parameter. Clearly, if we knew \mathbf{x}^* , we could compute $Q_{\mathbf{x}^*}(1 - \alpha)$ and calibrate our interval using $\psi_\alpha^2 := Q_{\mathbf{x}^*}(1 - \alpha) + s^2$. **But we don't!**
- **Core of the idea:** we can always obtain a $1 - \eta$ confidence set for \mathbf{x}^* by $f^{-1}(\Gamma_\eta(\mathbf{y})) := \left\{ \mathbf{x} \in \mathcal{X} : \|\mathbf{y} - \mathbf{K}\mathbf{x}\|_2^2 \leq \chi_{n,\eta}^2 \right\}$. We can then calibrate the interval by using $q_\gamma := \max_{\mathbf{x} \in f^{-1}(\Gamma_\eta(\mathbf{y}))} Q_x(1 - \gamma)$, such that $(1 - \eta)(1 - \gamma) = 1 - \alpha$. Uncertainty budget, trading off between confidence set (η) and quantile level (γ)
- This idea is similar to [Berger and Boos, 1994] and [Masserano et al., 2024], where tests involving nuisance parameters are controlled maximizing a p-value over a data-informed set.

We estimate q_γ using sampling and quantile regression

- Similar to the ideas present in [Dalmasso et al., 2020] [Dalmasso et al., 2022], [Masserano et al., 2023], Masserano et al., 2024]

adaOSB Algorithm

1. Let $\alpha, \gamma, \eta \in (0, 1)$ such that $\gamma(1 - \eta) + \eta = \alpha$, and $\Gamma_\eta(\mathbf{y})$ be a $1 - \eta$ confidence set containing $f(\mathbf{x}^*)$. Then $f^{-1}(\Gamma_\eta(\mathbf{y}))$ is a $1 - \eta$ confidence set for \mathbf{x}^* .
2. Generate samples $\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2, \dots, \tilde{\mathbf{x}}_M \sim \mathcal{U}(f^{-1}(\Gamma_\eta(\mathbf{y})))$. Task 1
3. Sample LLRs $\lambda_i \sim F_{\tilde{\mathbf{x}}_i}$ Task 2
4. Use generated data $\{(\tilde{\mathbf{x}}_i, \lambda_i)\}_{i=1}^M$ to fit a quantile regressor, $\hat{q}(\mathbf{x})$, and estimate $q_\gamma := \max_{\mathbf{x} \in f^{-1}(\Gamma_\eta(\mathbf{y})) \cap \mathcal{X}} \hat{q}(\mathbf{x})$
5. Compute $I(q_\gamma + s^2, \mathbf{y})$ with the guarantee that $\forall \mathbf{x}^* \in \mathcal{X}, \mathbb{P}_{\mathbf{y} \sim P_{\mathbf{x}^*}}(\varphi(\mathbf{x}^*) \in I(q_\gamma + s^2, \mathbf{y})) \geq 1 - \alpha$.

The **pros** and **cons** of this approach

- By finding $q_\gamma := \max_{\mathbf{x} \in f^{-1}(\Gamma_\eta(\mathbf{y}))} Q_{\mathbf{x}}(1 - \gamma)$, we avoid both
 1. Optimizing over a potentially unbounded space (e.g., $\mathcal{X} = \mathbb{R}_+^p$)
 2. Controlling for all $\mathbf{x} \in \mathcal{X}$ since we simply focus on the parameter values in $f^{-1}(\Gamma_\eta(\mathbf{y}))$ and adequately adjust the quantile we use.
- We **shift the complexity** to estimating q_γ :
 1. Sample generation is non-trivial - we develop two approaches for this.
 2. Estimating the max quantile via quantile regression

Implementation practicalities

Sampling $f^{-1}(\Gamma_\eta(\mathbf{y}))$

- For our examples, we focus on the scenario when $\mathcal{X} = \mathbb{R}_+^p$, $f(\mathbf{x}) = \mathbf{K}\mathbf{x}$, and $\varepsilon \sim N(\mathbf{0}, \Sigma)$, implying that we want to sample the intersection of the ellipsoid $\mathcal{E}(\mathbf{y}) := \{\mathbf{x} : \|\mathbf{y} - \mathbf{K}\mathbf{x}\|_2^2 \leq \chi_{n,\eta}^2\}$ and the non-negative orthant.
- **Two strategies**
 - Accept/Reject sampling uniformly from the pre-image ellipsoid is possible via [Voelker et al., 2017] but $\mathbb{P}(\mathbf{x}_i \in \mathbb{R}_+^p) \rightarrow 0$ as p gets large, and therefore becomes practically infeasible in higher dimensions (e.g., $p \geq 10$)
 - MCMC: Convex body [Smith, 1984] or polytope samplers [Chen et al., 2018]
 - We find a bounding polytope of $f^{-1}(\Gamma_\eta(\mathbf{y}))$ with hyperplanes defined by both the principal axes of the pre-image ellipsoid, the non-negativity constraints, and H additional randomly chosen hyperplanes.

Implementation practicalities

Quantile Regression

- Once we sample $\{(\tilde{\mathbf{x}}_i, \lambda_i)\}_{i=1}^M$, we perform quantile regression to learn $\hat{q}(\mathbf{x})$
- In principle, this regression can be done with any supervised learning algorithm using the pin-ball loss (e.g., [Meinshausen, 2006], [Takeuchi et al., 2006], [Dalmaso et al., 2020], [Dalmaso et al., 2021], [Masserano et al., 2023])
 - we use gradient-boosted regression since it has a clean implementation in sklearn.
- **Estimation of q_γ** : we sample in independent MCMC chain, $\{\bar{\mathbf{x}}_i\}_{i=1}^M$, and use the maximum out-of-sample predicted γ -quantile: $\hat{q}_\gamma := \max_{i \in [M]} \hat{q}(\bar{\mathbf{x}}_i)$
 - **Lemma 3.3** [Stanley et al., 2024]: \hat{q}_γ is a consistent estimator of q_γ .
 - **Theorem 1** [Dalmaso et al., 2021]: Quantile regression provides a consistent estimator of the quantile function.

Numerical Examples

Examples we consider

1. **Exposition** of method in simple 2d example [Tenorio et al., 2007] [Batlle et al. 2023]

$$\mathbf{y} = \mathbf{x} + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim N(\mathbf{0}, \mathbf{I}), \quad \varphi(\mathbf{x}) = x_1 - x_2, \quad \mathcal{X} = \mathbb{R}_+^2, \quad \mathbf{x}^* = (0.5 \quad 0.5)^T$$

2. **Valid Coverage** in a 3d scenario when OSB fails [Batlle et al. 2023]

$$\mathbf{y} = \mathbf{x} + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim N(\mathbf{0}, \mathbf{I}), \quad \varphi(\mathbf{x}) = x_1 + x_2 - x_3, \quad \mathcal{X} = \mathbb{R}_+^3, \quad \mathbf{x}^* = (0 \quad 0 \quad 1)^T$$

3. **Length Improvement** in a high dimensional ($p = 80$) scenario when OSB is empirically valid [Stanley et al. 2022] [Stanley et al. 2024a]

$$\mathbf{y} = \mathbf{K}\mathbf{x} + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim N(\mathbf{0}, \boldsymbol{\Sigma}), \quad \varphi(\mathbf{x}) = \mathbf{h}^T \mathbf{x}, \quad \mathcal{X} = \mathbb{R}_+^{80}, \quad \mathbf{x}^* \text{ defined mean bin counts}$$

We use $N = 1000$ samples to estimate interval coverage and length of OSB and adaOSB

Example 1: Exposition

2d

$$\mathbf{y} = \mathbf{x} + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim N(\mathbf{0}, \mathbf{I}), \quad \varphi(\mathbf{x}) = x_1 - x_2, \quad \mathcal{X} = \mathbb{R}_+^2, \quad \mathbf{x}^* = (0.5 \ 0.5)^T$$

- Since for any \mathbf{x} , we can efficiently estimate $Q_{\mathbf{x}}(1 - \gamma)$ in this example using Monte Carlo simulation, we do not use quantile regression, but rather use

$$\hat{q}_\gamma := \max_{i \in [M]} Q_{\tilde{\mathbf{x}}_i}(1 - \gamma)$$

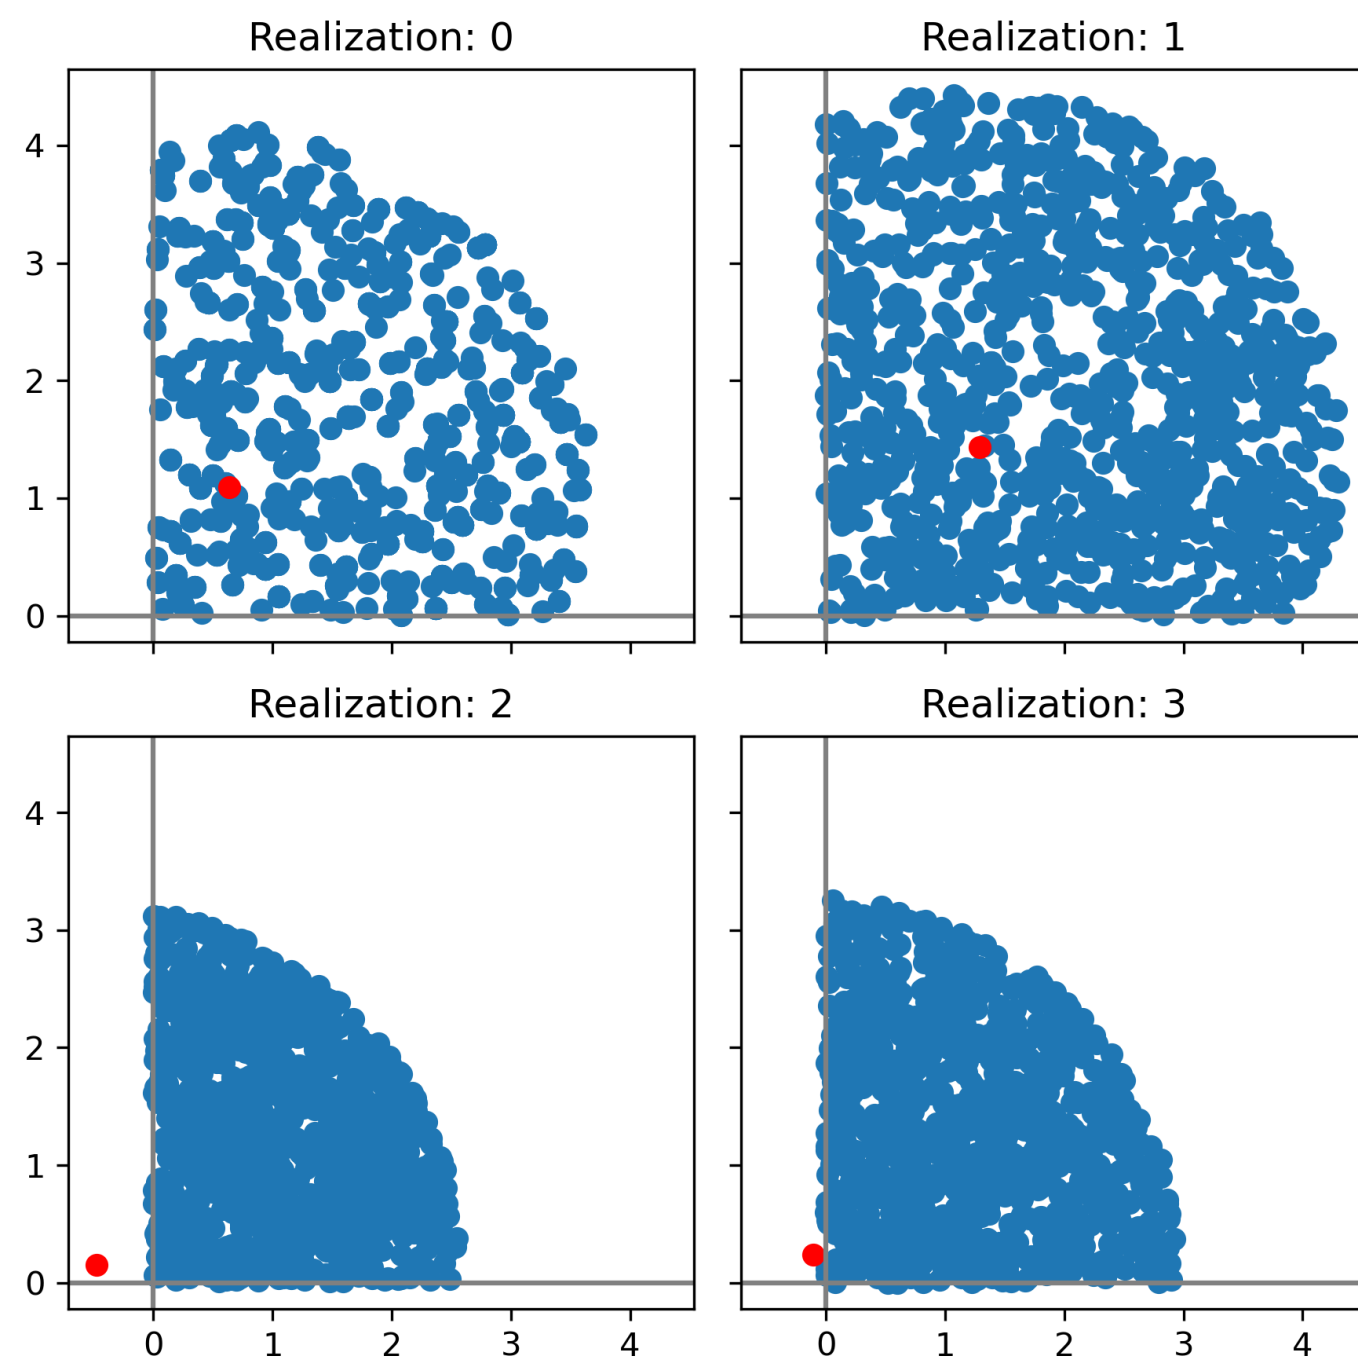
- We look to optimize a 68% interval ($\alpha = 0.32$). With $\eta := 0.01$, this implies that $\gamma = 0.3131$.
- Since $p = 2$, our accept/reject ellipsoid sampler is effective for sampling $\tilde{\mathbf{x}}_i \sim \mathcal{U}(f^{-1}(\Gamma_\eta(\mathbf{y})))$

Example 1: Exposition

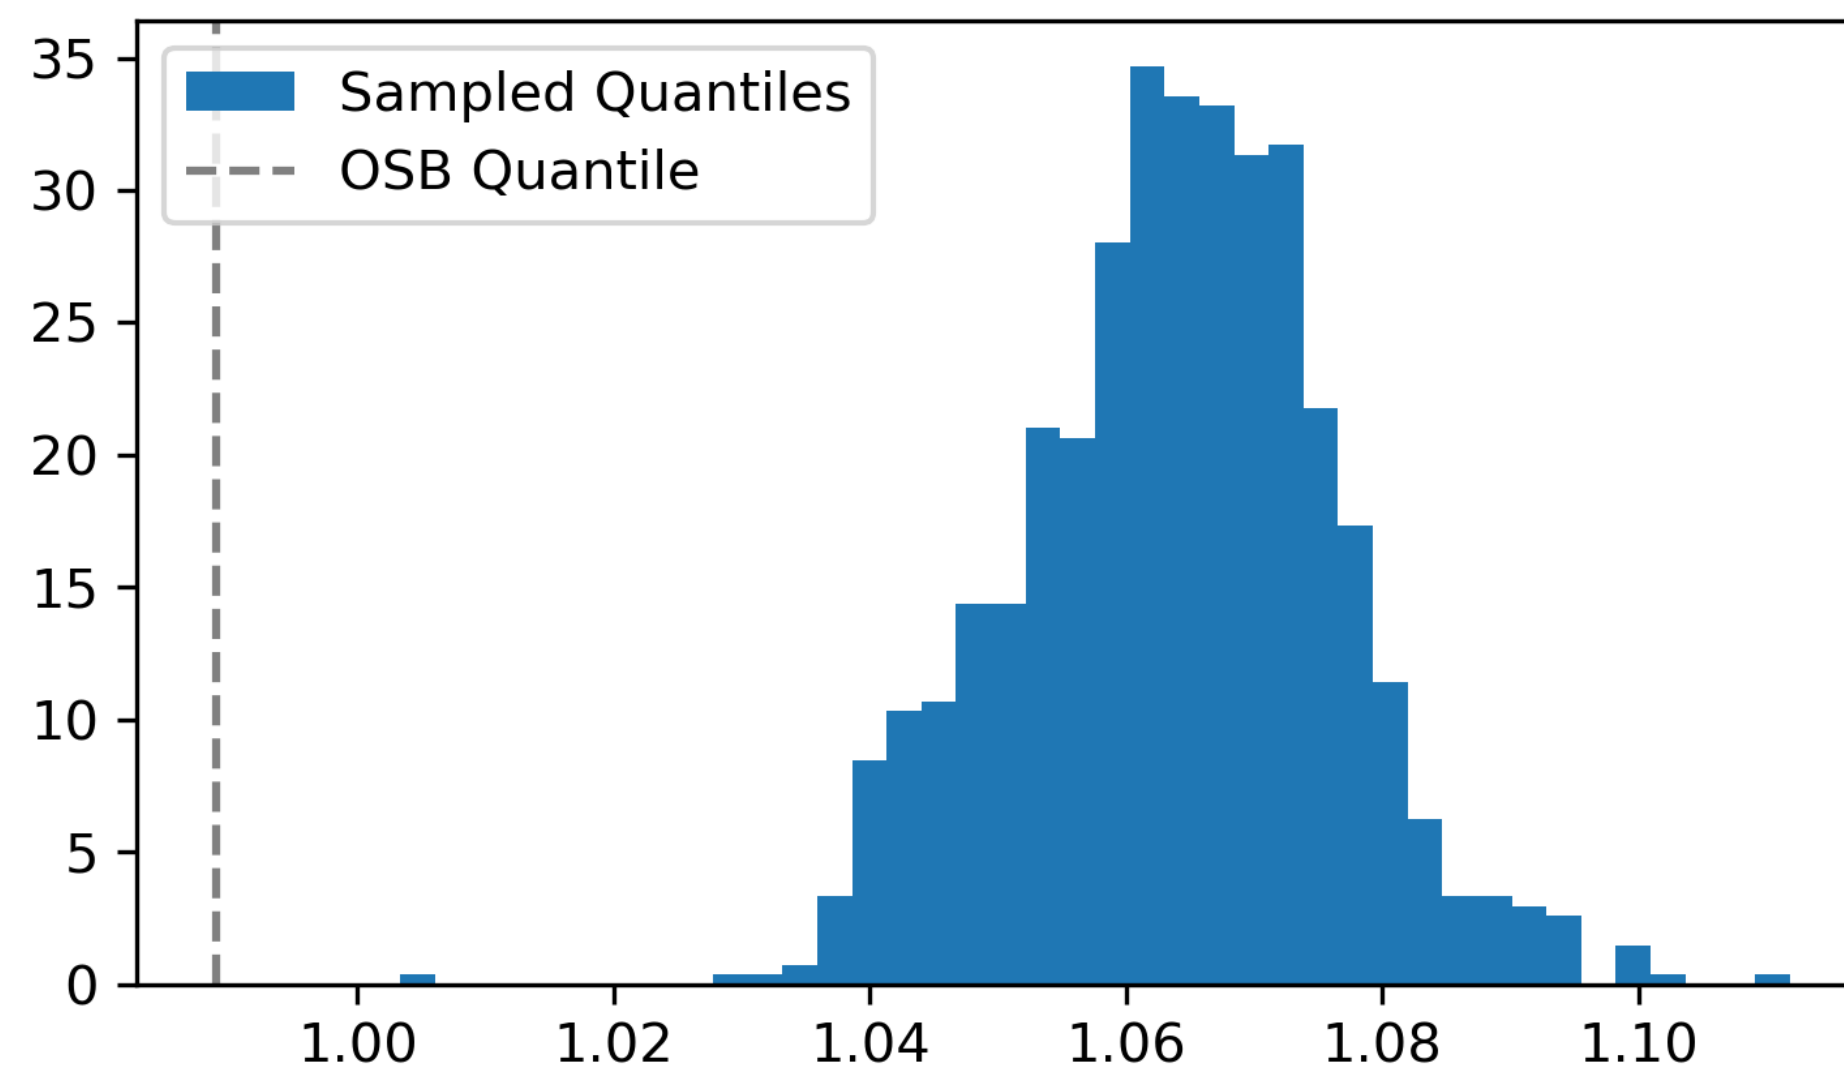
2d - We can see all the moving parts

$$\mathbf{y} = \mathbf{x} + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim N(\mathbf{0}, \mathbf{I}), \quad \varphi(\mathbf{x}) = x_1 - x_2, \quad \mathcal{X} = \mathbb{R}_+^2, \quad \mathbf{x}^* = (0.5 \quad 0.5)^T$$


Sampling $f^{-1}(\Gamma_\eta(\mathbf{y}))$



Observed distribution of \hat{q}_γ



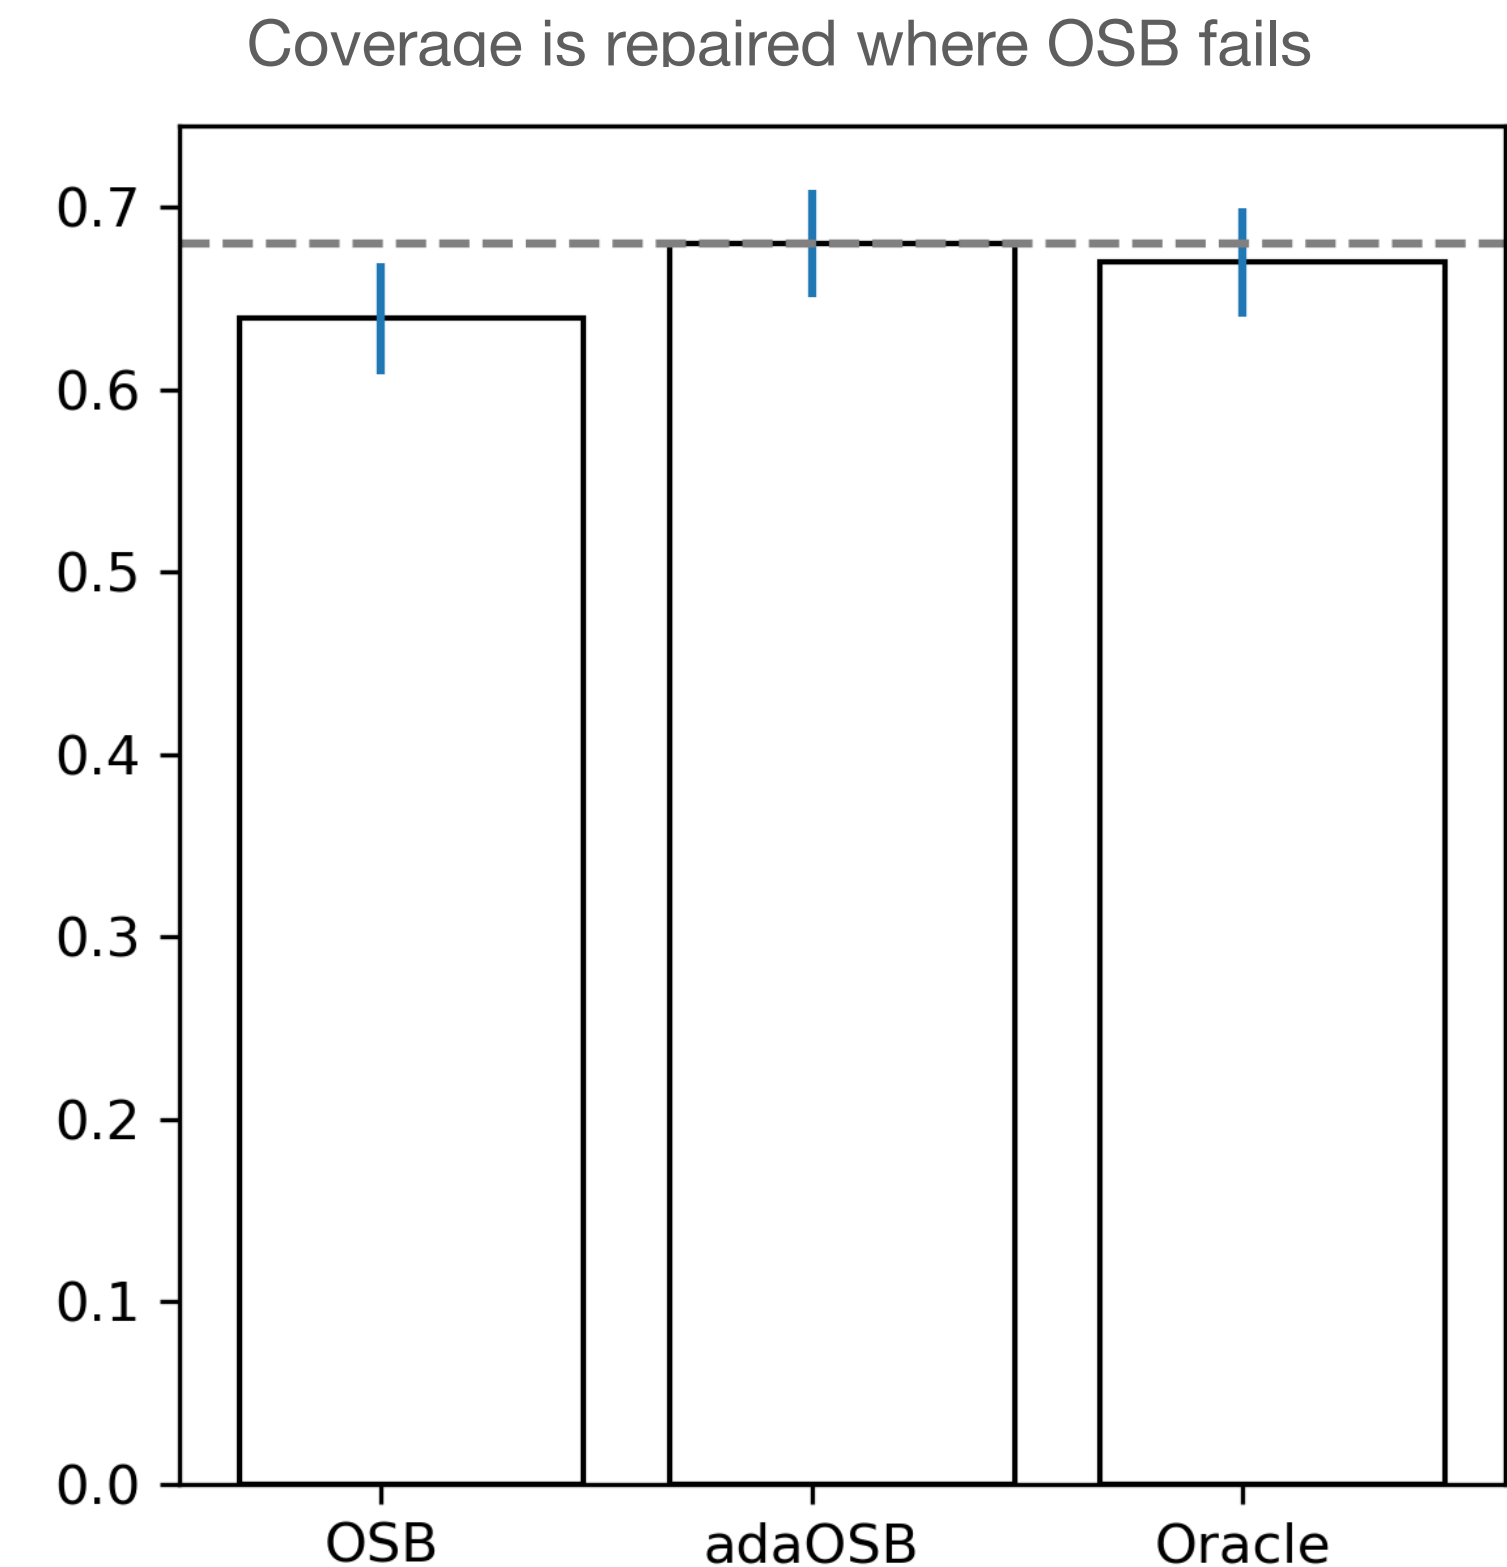
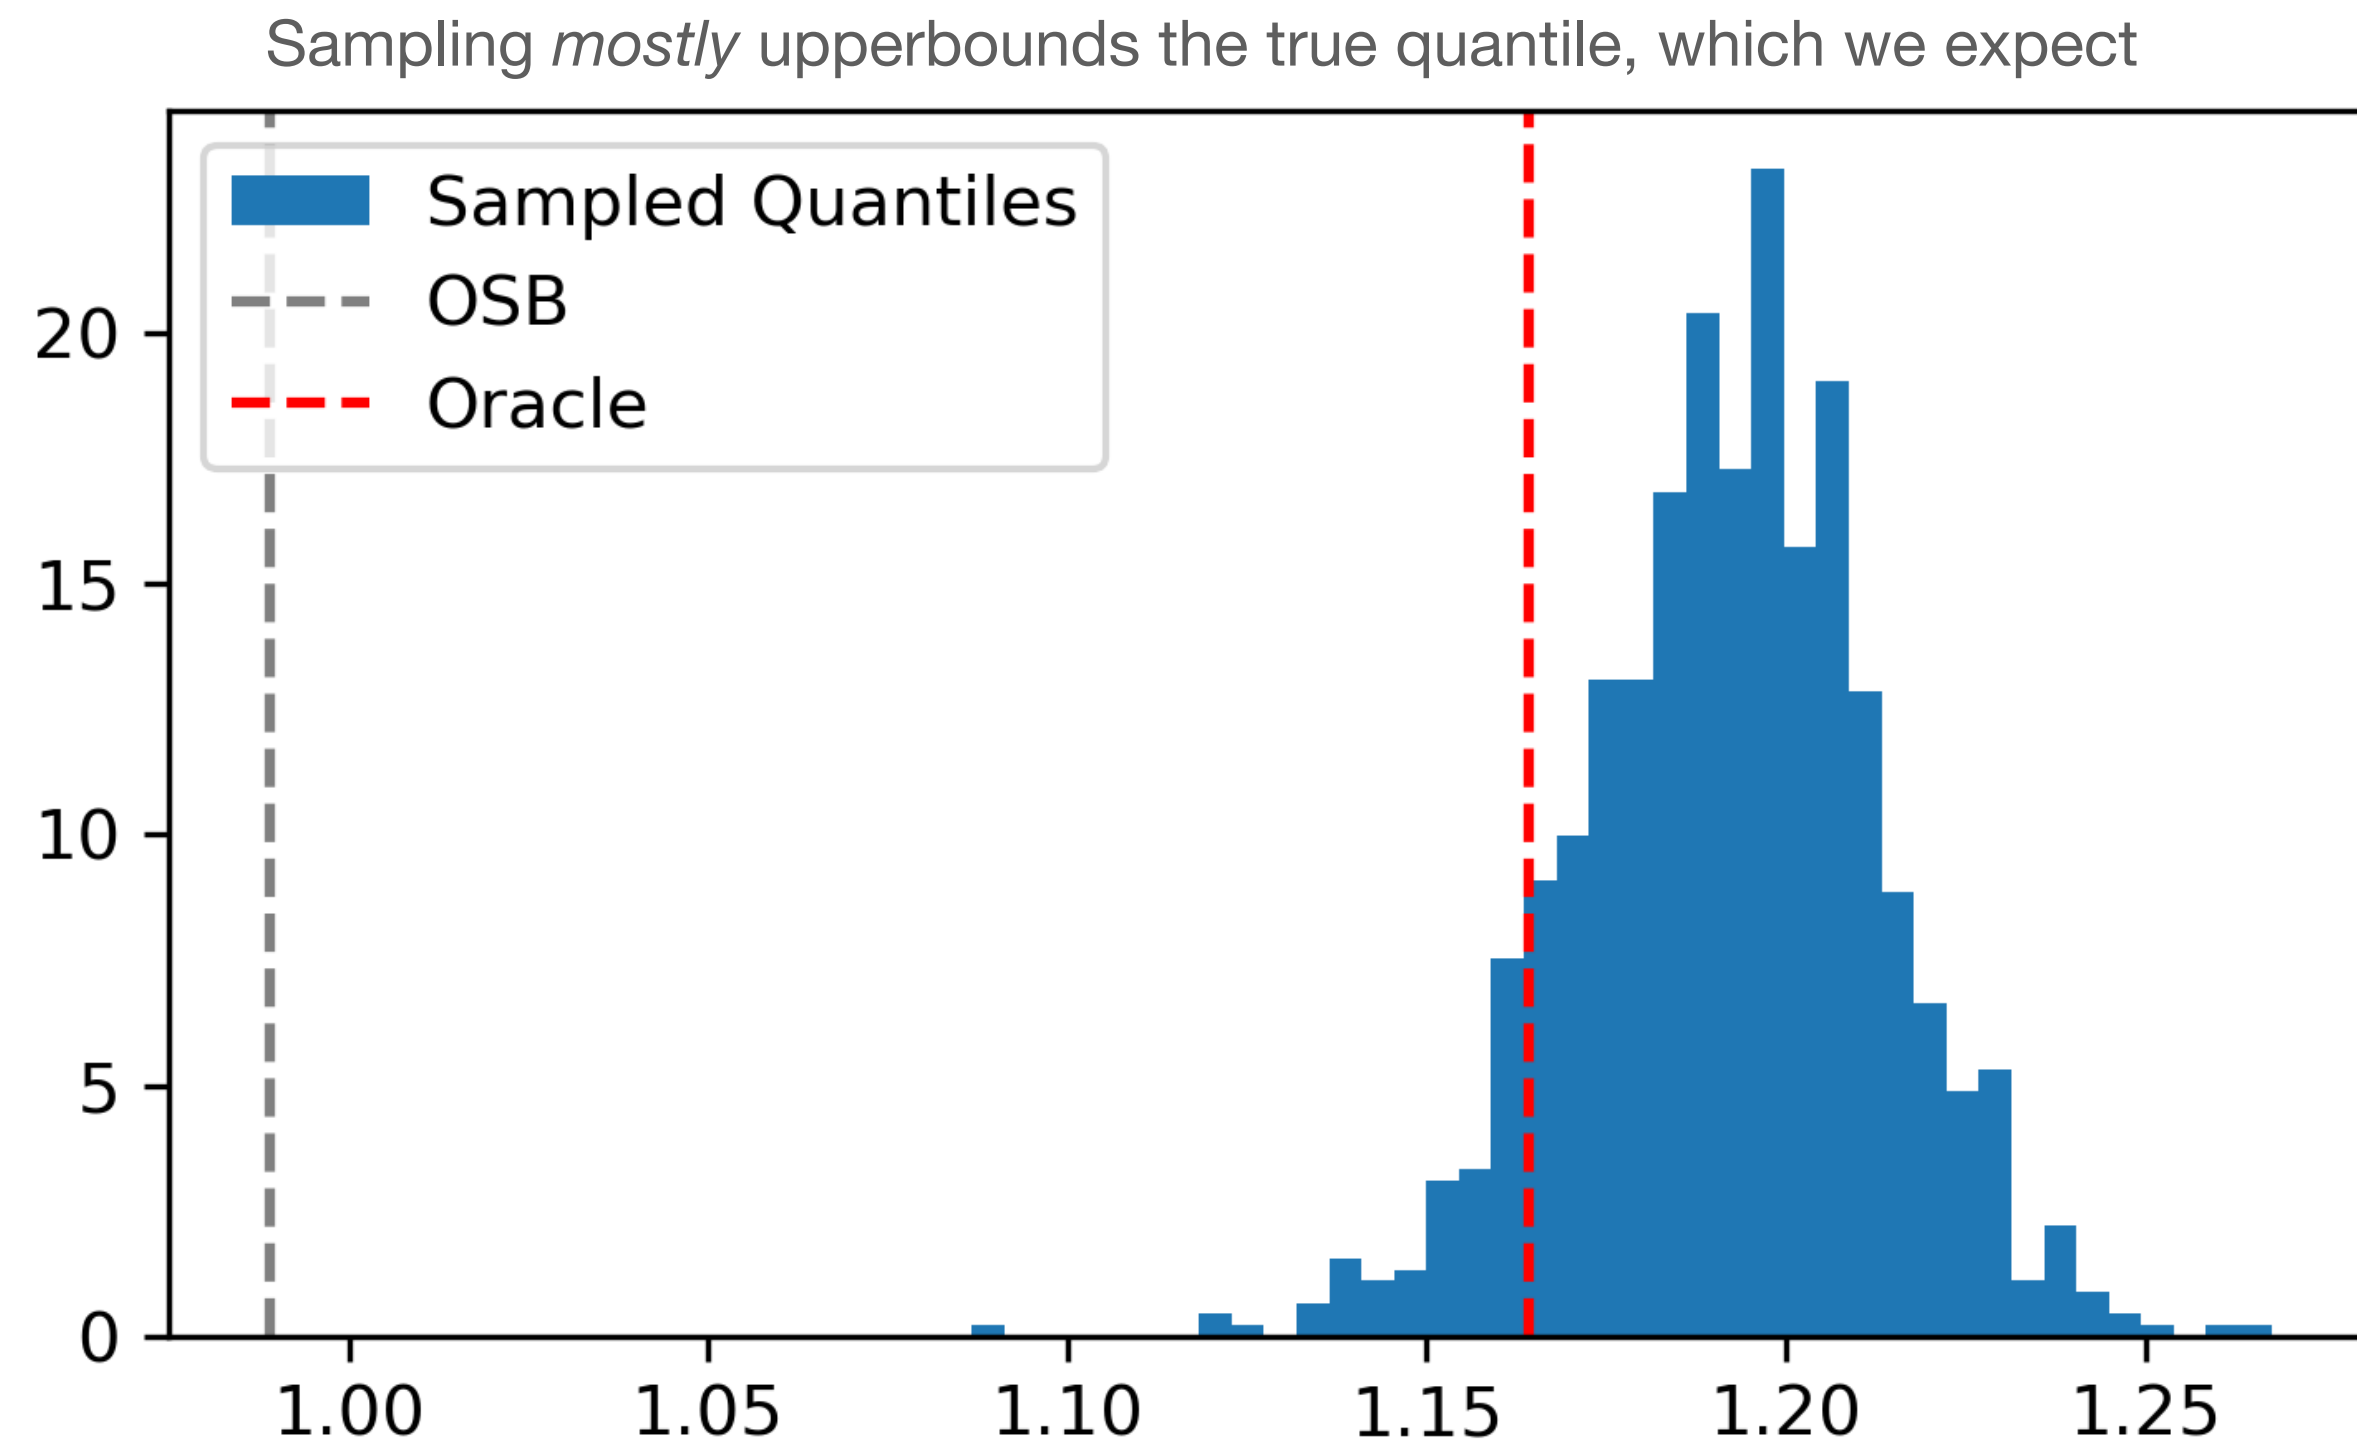
Estimated Interval Coverages



Example 2: Valid Coverage

3d - adaOSB adequately upper bounds true quantile and thus fixes coverage

$$\mathbf{y} = \mathbf{x} + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim N(\mathbf{0}, \mathbf{I}), \quad \varphi(\mathbf{x}) = x_1 + x_2 - x_3, \quad \mathcal{X} = \mathbb{R}_+^3, \quad \mathbf{x}^* = (0 \ 0 \ 1)^T$$




Example 3: Length Improvement

High dimension - Particle unfolding simulation where adaOSB shows a dramatic length improvement


$$\mathbf{y} = \mathbf{K}\mathbf{x} + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim N(\mathbf{0}, \boldsymbol{\Sigma}), \quad \varphi(\mathbf{x}) = \mathbf{h}^T \mathbf{x}, \quad \mathcal{X} = \mathbb{R}_+^{80}, \quad \mathbf{x}^* \text{ defined mean bin counts}$$

- High dimensions necessitate MCMC polytope sampler and quantile regression


Distributions of MCMC/QR derived \hat{q}_γ



adaOSB has a **coverage** advantage



adaOSB has a clear **length** advantage



Recap and conclusions

1. Building on the work of [Batlle et al. 2023], we presented a method to set ψ_α^2 in a data-dependent way to achieve interval coverage and improve interval length relative to OSB.
 - **Take-away:** our method is the first computationally feasible approach to properly calibrate these optimization-based intervals.
 - **Key Steps:** using an uncertainty budget to bound the set of feasible parameter values, sampling the pre-image confidence set, estimating the max quantile.
2. We explored three numerical studies to demonstrate the method and its advantages.
 - **Take-away:** our method provides coverage in low dimensional ($p = 3$) example where OSB does not, and improves interval length in a scenario where OSB empirically over-covers ($p = 80$).

Thank You!

Please let me know if you have any follow up questions:
mcstanle@andrew.cmu.edu

References

- (1) Batlle, P., Stanley, M., Patil, P., Kuusela, M., Owhadi, H. (2023). Confidence intervals for functionals in constrained inverse problems: An optimization-centric perspective with frequentist guarantees. *Arxiv*.
- (2) Roger L. Berger and Dennis D. Boos. (1994). *P Values Maximized Over a Confidence Set for the Nuisance Parameter*. *Journal of the American Statistical Association*.
- (3) Chen, Y., Dwivedi, R., Wainwright, M., Yu, B. (2018). Fast MCMC Sampling Algorithms on Polytopes. *Journal of Machine Learning Research*.
- (4) Dalmaso, N., Izbiki, R., Lee, A. (2020) Confidence and Hypothesis Testing in a Likelihood-Free Inference Setting. *International Conference on Machine Learning*.
- (5) Dalmaso, N., Masserano, L., Zhao, D., Izbiki, R., Lee, A. (2021). Likelihood-Free Frequentist Inference: Bridging Classical Statistics and Machine Learning For Reliable Simulator-Based Inference. *Arxiv*.
- (6) Han, T., Low, M., Wu, T. A. (2023). Quantum Entanglement and Bell Inequality Violation in Semi-Leptonic Top Decays. *Arxiv*.
- (7) Kuusela, M. (2016). Uncertainty quantification in unfolding elementary particle spectra at the Large Hadron Collider. *PhD Thesis*.
- (8) Masserano, L., Dorigo, T., Izbiki, R., Kuusela, M., Lee, A. (2023). Simulator-based Inference with WALDO: Confidence Regions by Leveraging Prediction Algorithms and Posterior Estimators for Inverse Problems. *International Conference on Artificial Intelligence and Statistics*.
- (9) Masserano, L., Shen, A., Doro, M., Dorigo, T., Izbiki, T., Lee, A. (2024). Classification Under Nuisance Parameters and Generalized Label Shift in Likelihood-Free Inference. *Arxiv*.
- (10) Meinshausen, N. (2006). Quantile Regression Forests. *Journal of Machine Learning Research*.
- (11) Patil, P., Kuusela, M., Hobbs, J. (2022). Objective frequentist uncertainty quantification for atmospheric CO2 retrievals. *SIAM / ASA Journal on Uncertainty Quantification*.
- (12) B. W. Rust and W. R. Burrus (1972). Mathematical programming and the numerical solution of linear equations. *American Elsevier Publishing Company*.
- (13) B. W. Rust and D. P. O'Leary (1994). Confidence intervals for discrete approximations to ill-posed problems. *Journal of Computational and Graphical Statistics*. 3.
- (14) Smith, R.L. (1984). Efficient Monte Carlo Procedures for Generating Points Uniformly Distributed Over Bounded Regions. *Operations Research*.
- (15) Stanley, M., Patil, P., Kuusela, M. (2022). Uncertainty quantification for wide-bin unfolding: one-at-a-time strict bounds and prior-optimized confidence intervals. *Journal of Instrumentation*.
- (16) Stanley, M., Batlle, P., Kuusela, M., Owhadi, H. (2024a). AdaOSB: adaptive one-at-a-time strict bounds for functionals in constrained inverse problems. *In progress*.**
- (17) Stanley, M., Kuusela, M., Byrne, B., Liu, J. (2024b). Optimization-based confidence intervals for regional carbon flux inversion uncertainty quantification. *In progress*.
- (18) Stark, P. (1992). Inference in Infinite-Dimensional Inverse Problems: Discretization and Duality. *Journal of Geophysical Research*.
- (19) Takeuchi, I., Q. V., Sears, T. D., and Smola, A. J. (2006). Nonparametric quantile estimation. *Journal of Machine Learning*, 7:1231-1264.

Appendix

Developing the test inversion formalism in this setting provides a new perspective

Key set definitions $\mathcal{X} \subset \mathbb{R}^p$ $\Phi_\mu := \{x : \varphi(x) = \mu\} \subset \mathbb{R}^p$

Fundamental HT $H_0 : x^* \in \Phi_\mu \cap \mathcal{X}$ versus $H_1 : x^* \in \mathcal{X} \setminus \Phi_\mu$

Test Statistic (LLR) $\lambda(\mu, \mathbf{y}) := -2 \log \Lambda(\mu, \mathbf{y}) = -2 \left(\sup_{x \in \Phi_\mu \cap \mathcal{X}} \ell_x(\mathbf{y}) - \sup_{x \in \mathcal{X}} \ell_x(\mathbf{y}) \right)$
 $= \inf_{x \in \Phi_\mu \cap \mathcal{X}} -2\ell_x(\mathbf{y}) - \inf_{x \in \mathcal{X}} -2\ell_x(\mathbf{y})$

Level α test $\sup_{x \in \Phi_\mu \cap \mathcal{X}} \mathbb{P}_{\lambda \sim F_x} (\lambda > q_\alpha) \leq \alpha$ Test T_μ is a level- α test

Let $Q_x : [0, 1] \rightarrow \mathbb{R}$ be the quantile function of $\lambda(\mu, \mathbf{y})$ at x . Using $Q_x(1 - \alpha)$ produces a *level- α* test.

Ellipsoid Sampler

Uniform sampling in p -ball + Accept/reject

- [Voelker et al., 2017] presented and proved an interesting and efficient algorithm to sample uniformly at random from the p -ball. First, sample uniformly from the $(p + 1)$ -sphere (possible with Gaussian RNG) followed by dropping any two coordinates.
 - We refer to a sample drawn from the p -ball via “Voelker-Gosmann-Stewart” (VGS) by $\mathbf{x} \sim VGS(p)$
- Consider an ellipsoid defined by $\mathcal{E}(r) := \{\mathbf{x} : \mathbf{x}^T \mathbf{A} \mathbf{x} \leq r\}$ and let $\mathbf{P} \mathbf{\Omega}^2 \mathbf{P}^T$ be the eigendecomposition of PSD \mathbf{A} .
- If $\mathbf{x} \sim VGS(p)$, then $\mathbf{y} := \sqrt{\chi_{n,\eta}^2} \mathbf{P} \mathbf{\Omega} \mathbf{x}$ is sampled uniformly at random from $\mathcal{E}(\chi_{n,\eta}^2)$
- To incorporate constraints, simple reject \mathbf{y} if $\mathbf{y} \notin \mathcal{X}$
- NOTE: this approach works well in low dimensions and when $f(\mathbf{x}) = \mathbf{K} \mathbf{x}$, where \mathbf{K} is full column rank.

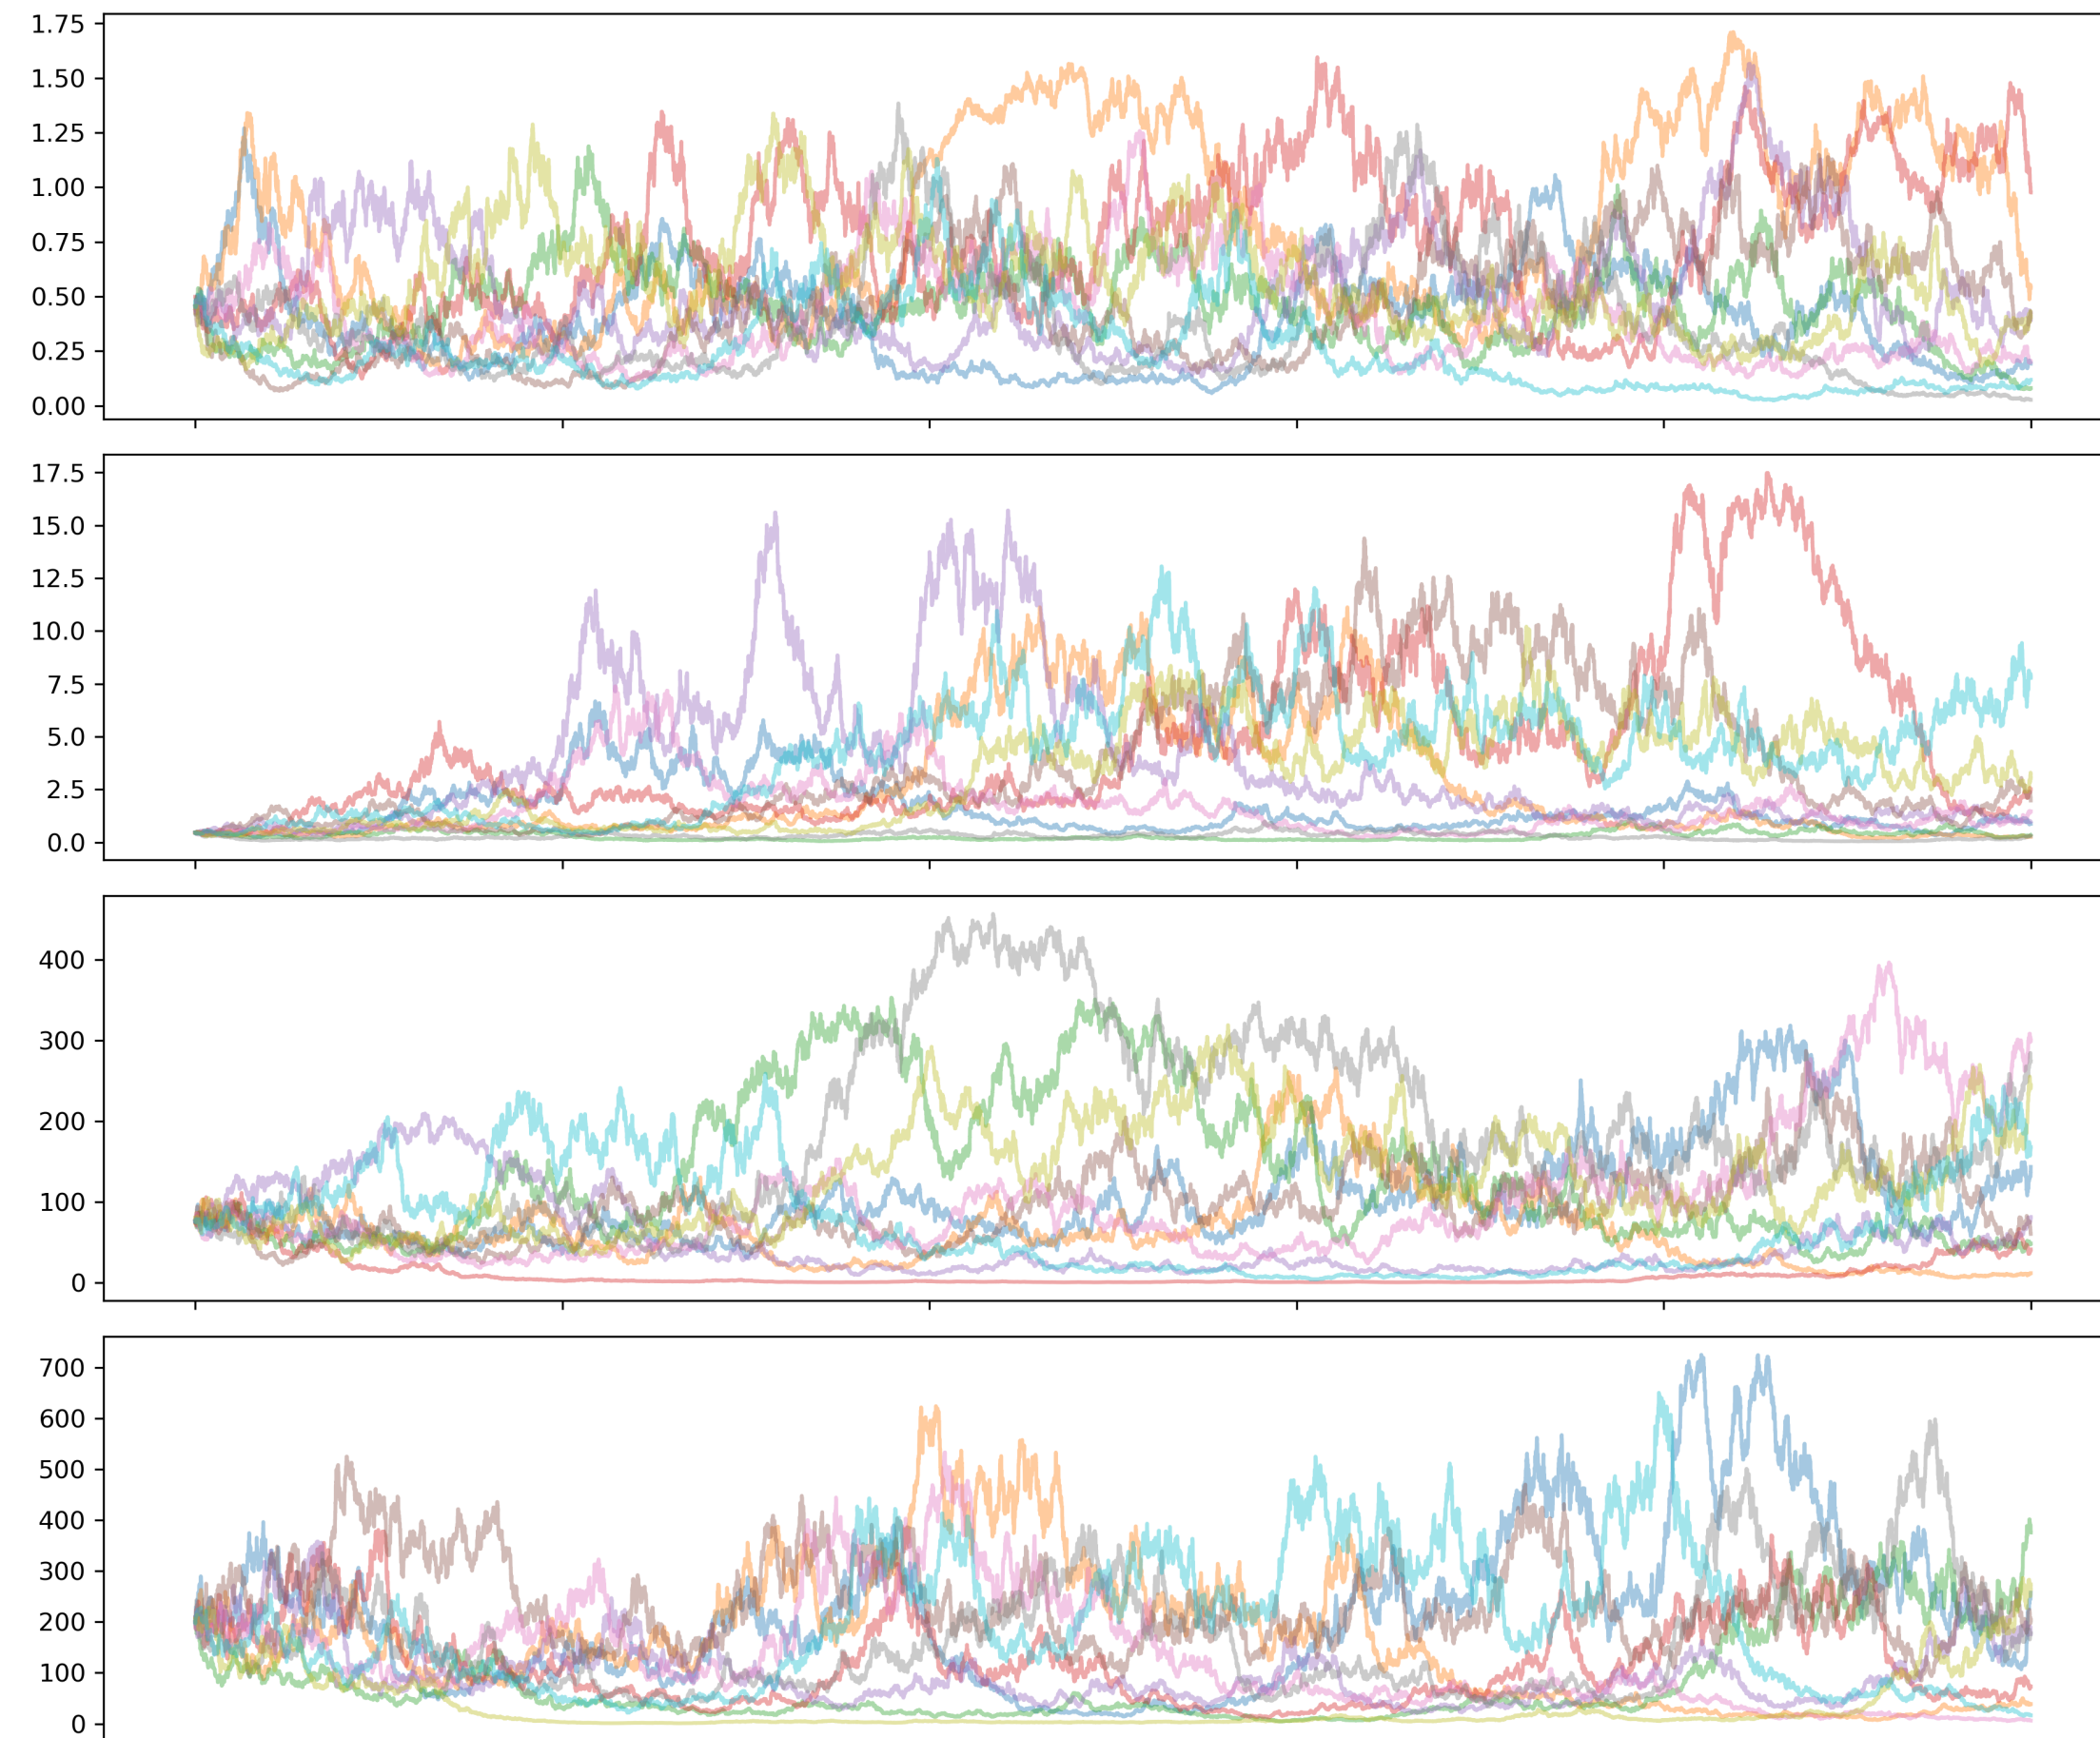
MCMC Polytope Sampler

Implementation details and considerations

- We construct a bounding polytope for $f^{-1}(\Gamma_\eta(\mathbf{y}))$ using the principal axes of the confidence set ellipsoid ($2p$), the hyper-rectangle defined by the non-negativity constraints ($2p$) and 200 additional randomly chosen hyperplanes.
- We use the Vaidya sampler detailed in [Chen et al., 2018], where the uniform distribution over the defined polytope is the Markov chain's stationary distribution
- Since this sampling is an MCMC algorithm, we consider a few different convergence plots to assess sufficient mixing:
 - Trace plots of individual parameters
 - Ensembles of max predicted quantiles for both **fixed data set** size and **cumulative**
 - **Fixed** allows for us to get a sense of the Markov chain convergence
 - **Cumulative** allows us to assess the stability the max predicted quantile

MCMC Polytope Sampler (con't)

Parameter Trace Plots




Four arbitrarily chosen parameter trace plots show nice mixing

MCMC Polytope Sampler (con't)

Fixed Max-q trace plots

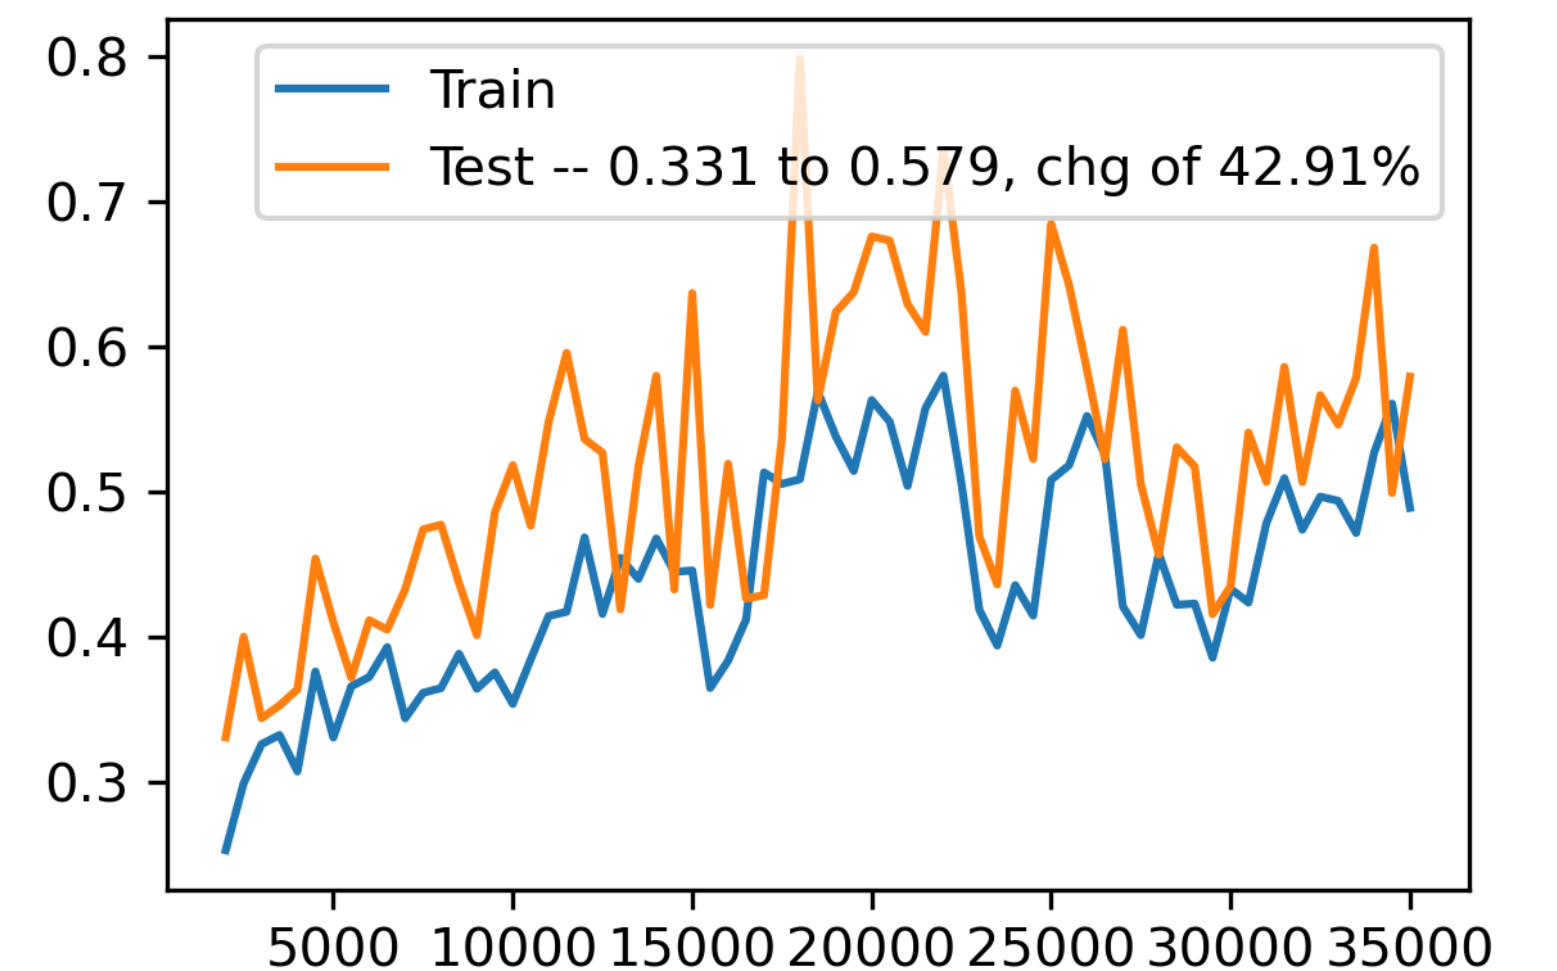
The ensemble of Max predicted quantiles with dataset size fixed at 2k



Ensemble mean stabilizes after ~15k iterations




Ensemble width stabilizes after ~15k iterations




MCMC Polytope Sampler (con't)

Cumulative Max-q trace plots


The ensemble of Max predicted quantiles with dataset size fixed at 2k



Ensemble mean stabilizes after ~10k iterations




Ensemble width decreases



2d Exposition example

Additional figures and details

Realization 0 with pax estimated quantile



Monte Carlo Sampling to estimate $Q_x(1 - \gamma)$

1. Generate an ensemble of samples, $\mathbf{y}_i = \mathbf{x} + \boldsymbol{\varepsilon}_i$ and therefore LLR samples, $\lambda(\mathbf{h}^T \mathbf{x}, \mathbf{y}_i)$.
2. From our generated ensemble, we can simply use the $(1 - \gamma)$ percentile estimator.

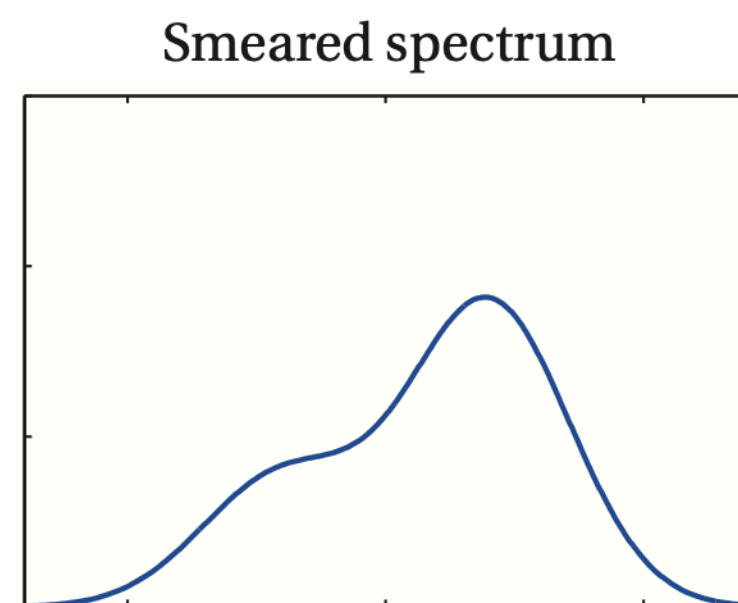
More on particle unfolding

The data generating process for our histogram is

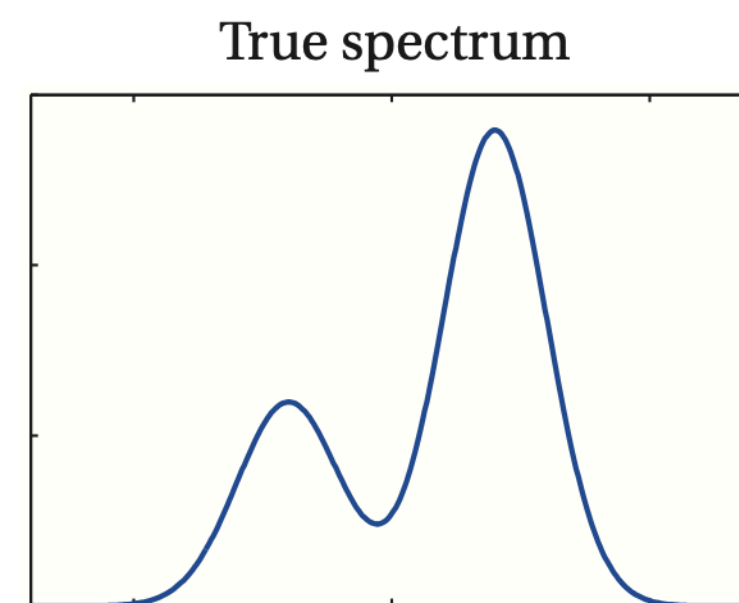
$$\mathbf{y} \sim \text{Poisson}(\mathbf{K}\boldsymbol{\lambda}),$$

which we approximate by

$$\mathbf{y} \sim N(\mathbf{K}\hat{\boldsymbol{\lambda}}, \boldsymbol{\Sigma}), \quad \Sigma_{ii} = (K\hat{\lambda})_i, \quad \forall i.$$



Folding
←
Unfolding
→



For more information, see [Kuusela, 2016] and [Stanley et al., 2022]