

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 WEIGHTED CONFORMAL PREDICTION FOR TIME- DEPENDENT PDES

Anonymous authors

Paper under double-blind review

ABSTRACT

Uncertainty quantification is crucial in scientific machine learning, where models inform safety-critical tasks such as flood forecasting, aerodynamic optimization, and financial risk management. Conformal prediction provides distribution-free coverage guarantees, but in time-dependent settings common to physics and engineering, these guarantees can break down, leading to systematic undercoverage. We study this problem in the context of surrogate models for time-dependent physical systems described by partial differential equations (PDEs). We prove that in a function space setting, distributions at arbitrarily close times can be mutually singular, making exact coverage guarantees impossible. As a solution, we facilitate weighted conformal prediction for a broad class of PDE problems arising from discretized models and validate these results in experiments. While prior work often sidesteps time dependence—by assuming exchangeability, focusing on short horizons, or ignoring long-term deployment—we address it directly by providing exact coverage guarantees through reweighting calibration scores.

1 INTRODUCTION

Many problems in physics and engineering, such as weather prediction, aerodynamics, and financial modeling, are governed by partial differential equations (PDEs). Classical numerical solvers are accurate but computationally expensive, scaling poorly with dimensionality or repeated simulations. AI-based surrogate models have emerged as a promising alternative, providing fast approximations of PDE solutions. Prominent examples include physics-informed neural networks (PINNs) (Raissi et al., 2019), DeepONets (Lu et al., 2021), and neural operators (Anandkumar et al., 2019; Li et al., 2021). Most notably, neural operators have demonstrated remarkable success in generalizing across different discretizations, geometries, and boundary conditions.

Despite these advances, surrogate models still lack principled mechanisms for uncertainty quantification. This limitation is critical, since scientific and engineering decisions often depend on reliable confidence assessments of model outputs. Conformal prediction (CP) Vovk et al. (2022) provides a principled framework, producing distribution-free uncertainty sets with guaranteed marginal coverage. These guarantees, however, rely on exchangeability between calibration and test samples—a condition that is frequently violated in time-dependent PDEs.

Non-Stationarity in Time-Dependent PDEs. Let u_t denote the solution of a time-dependent PDE at time point t . In practice, we are interested in predicting $u_{t+\delta}$ for several time steps $\delta > 0$, beyond the available training and calibration data. Unless u_t is a stationary process, test samples follow a different distribution than observed calibration samples, breaking the exchangeability assumption required by conformal prediction.

This type of non-stationarity is ubiquitous: sudden shocks (e.g., stock market crashes), long-term structural changes (e.g., climate trends), and limited development windows (e.g., laboratory testing) all produce systematic shifts in the data distribution (see figure 1). Even for simple PDEs, the marginal distribution of u_t may drift continuously in t and diverge arbitrarily as $t \rightarrow \infty$.

Implications for Conformal Prediction. The consequence is that conformal intervals calibrated at time t may undercover at future times $t + \delta$. Figure 2 illustrates this behavior on the backward heat equation. In the top row, calibration at time step δ still yields valid coverage at 3δ . In contrast,

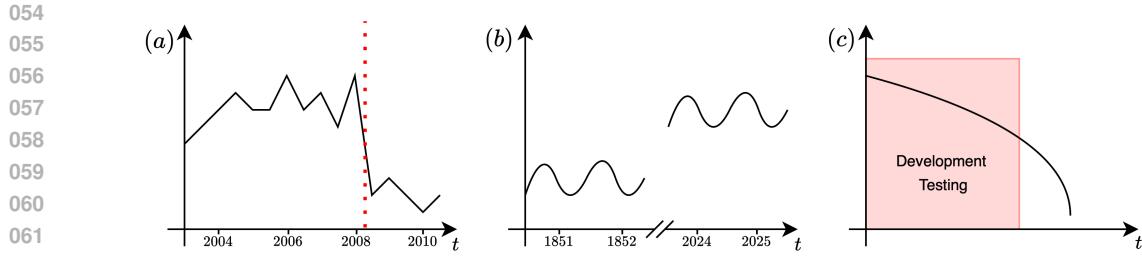


Figure 1: Examples of different types of non-stationarity that arise in time-dependent systems. (a) Sudden distribution shift due to external shocks, illustrated by a stock market crash. (b) Long-term structural changes, as in climate time series. (c) Limited observation window during system development, where calibration and testing occur on different parts of the trajectory.

calibration at step 4δ (bottom row) leads to undercoverage already at the first prediction time step, because the PDE grows progressively unstable. This demonstrates that simply tuning the discretization step does not resolve the problem: coverage degradation is inherent to the temporal drift.

A concrete example arises in weather forecasting, where calibration on short-term simulations may produce intervals that appear reliable but fail to capture rare extreme events at later horizons. In such cases, nominal 90% coverage can collapse well below the target, producing forecasts that systematically underestimate risk.

Consequently, CP coverage guarantees do not hold in time-dependent PDEs. While there are first methods to sidestep the non-exchangeability (see section 2), all of these come with limiting assumptions that prohibit broad applicability.

Our Contributions. In this work, we address this gap by studying CP for time-dependent surrogate models of PDEs, providing the following contributions:

1. We analyze the function-space formulation of the learning problem and show that even in simple settings, such as the heat equation, the total variation (TV) distance is maximal for any time distance. This shows that a pure function-space perspective, as often used in the neural operator literature, is unsuitable for the non-exchangeable CP framework.
2. For a broad class of PDEs, we derive explicit densities for the discretized solutions over time, facilitating the use of *weighted conformal prediction*. This enables exact coverage guarantees for PDEs without limiting assumptions on their time-dependent behavior.
3. We empirically validate our method on several time-dependent PDEs and compare it to alternative CP approaches (which assume exchangeability or local exchangeability). We show that these limiting assumptions on the time dynamics indeed lead to undercoverage, and that our approach is the only method providing reliable coverage over time.

The paper is structured as follows. In section 2, we review related work. Section 3 provides background on CP, PDEs, and surrogate models. In section 4, we formalize the problem setting, present our result on function spaces, and our weighted CP framework. Section 5 presents empirical results demonstrating the effectiveness of our approach and section 6 concludes.

2 RELATED WORK

Trajectory-Based Exchangeability. The most straightforward option to bypass the exchangeability issue is to treat entire trajectories as the exchangeable units. Moya et al. (2025) use DeepONets to predict full solution trajectories, calibrating CP on trajectory-level samples. This avoids assumptions on exchangeability within the calibrated time horizon, but does not address potential distribution shifts beyond this horizon, e.g. in a potential model deployment. Gray et al. (2025) follow the same strategy, though their method applies to arbitrary surrogate models beyond neural operators. Gopakumar et al. (2025) also adopt trajectory-level calibration, but focus on conformal sets for deviations between surrogates and the governing PDE operator, rather than for the solution itself.

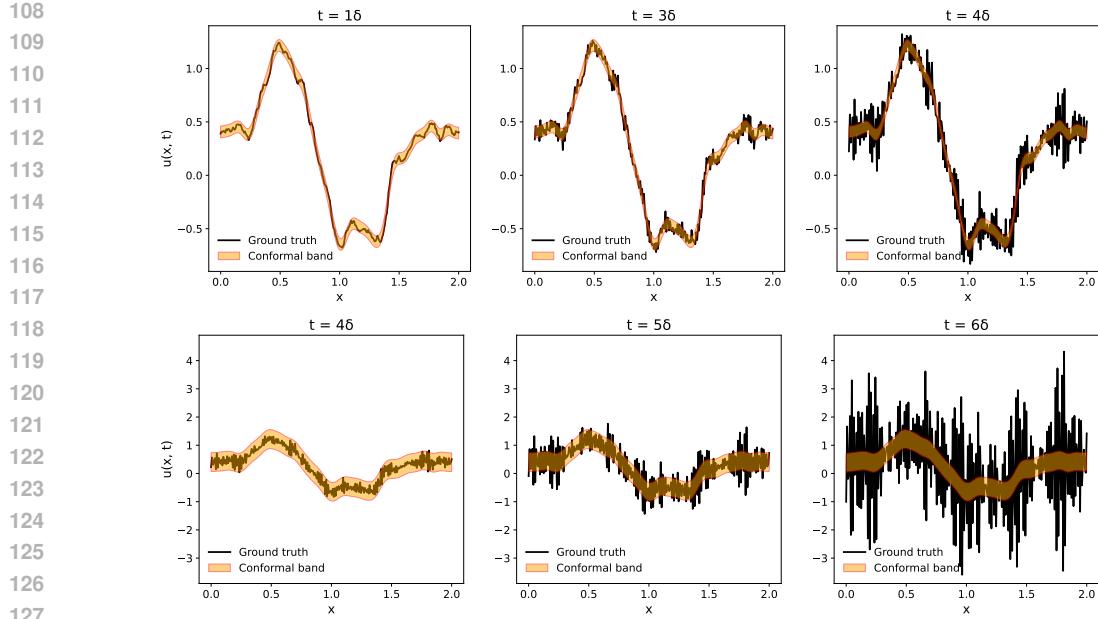


Figure 2: CP bands when calibrating at different time resolutions of the backward heat equation. Each figure shows the solution $u(x, t)$ and CP bands over the spatial domain x at one time point. (Top) Calibration on data from time step δ : prediction bands remain valid even at step 3δ . (Bottom) Calibration on data from time step 4δ : undercoverage occurs already after one time step.

Relaxed Exchangeability Assumptions. Motivated by data scarcity, Gopakumar et al. (2024) go beyond trajectory-based exchangeability and construct calibration samples by slicing long time series into shorter segments. This construction implicitly assumes that the time series is stationary across segments, which may hold approximately in periodic systems (e.g., weather data), but fails in general dynamical systems with non-periodic trends or sustained drifts.

Harris & Liu (2025) take a different approach with their Local Spectral Conformal Inference (LSCI) method, designed for neural operators. Instead of global exchangeability, they derive conformal bands with guarantees under *local exchangeability*, i.e., that points close in time are nearly exchangeable. When local exchangeability holds, LSCI provides the first principled way for time-adaptive prediction sets with coverage guarantees for neural operators. However, validating this assumption in practice is usually not feasible (see appendix A.1). Therefore, Harris & Liu (2025) *assume* local exchangeability in their experiments by taking very small time steps. In Figure 2, however, we see one example where a calibration on time step 1 leads to good empirical coverage for three further time steps, but when calibrating again at a later time, coverage already drops after one time step because the solution gets exponentially noisier. Thus, tuning the step size at calibration does not ensure local exchangeability at test time.

Time Series Conformal Prediction Time-series conformal prediction is an active research area. Most existing methods provide only asymptotic guarantees (e.g., Gibbs & Candès (2021); An-gelopoulos et al. (2023); Xu & Xie (2023)), meaning they achieve coverage on average over infinite time steps or only in the limit as the time horizon grows. Another line of work seeks subclasses of time series with structural properties that enable per-time-step coverage. For example, Oliveira et al. (2022) shows coverage for stationary, β -mixing processes. Our approach falls into this latter category, and to the best of our knowledge, no prior work leverages PDE-specific properties to obtain conformal coverage.

162

3 BACKGROUND

164

3.1 CONFORMAL PREDICTION

166 Conformal prediction (CP) is a framework for constructing prediction sets with marginal finite-
 167 sample coverage guarantees Vovk et al. (2022). In the standard split setting, a model is trained on
 168 $\mathcal{D}_{\text{train}}$ and calibrated on \mathcal{D}_{cal} , yielding a set-valued predictor \mathcal{C} such that, for a test sample (\mathbf{x}, \mathbf{y}) ,

$$169 \quad \mathcal{P}(\mathbf{y} \in \mathcal{C}(\mathbf{x})) \geq 1 - \alpha \\ 170$$

171 at coverage level $1 - \alpha$. This guarantee relies on **exchangeability** of calibration and test samples—that is, their joint distribution is invariant under permutations. When exchangeability is vio-
 172 lated, coverage may fail.

174 **Conformal Prediction Beyond Exchangeability.** When calibration and test distributions differ
 175 but are related by a likelihood ratio, weighted CP provides a natural extension Vovk et al. (2022);
 176 Barber et al. (2023). In this setting, calibration samples are reweighted by
 177

$$178 \quad w_i \propto \frac{p_{\text{test}}(\mathbf{x}_i)}{p_{\text{cal}}(\mathbf{x}_i)}, \quad \sum_i w_i = 1, \\ 179 \\ 180$$

181 so that the conformal quantile is computed with respect to these weights. Here, the index i ranges
 182 over all calibration data points and the target test point. If the density ratio is known or can be
 183 estimated, weighted CP can restore exact coverage in covariate-shift settings. In our PDE setup,
 184 the linear–Gaussian structure allows us to compute these ratios in closed form, enabling precise
 185 conformal bands (see section 4.4).

186 In case a closed-form evaluation of the weights is not possible, Barber et al. (2023) provide correc-
 187 tions for the conformal guarantees based on the TV distance¹ between calibration and test distribu-
 188 tion, that hold even in the general case of non-exchangeability:

$$189 \quad \mathcal{P}(\mathbf{y} \in \mathcal{C}(\mathbf{x})) \geq 1 - \alpha - \sum_{i=1}^n w_i d_{\text{TV}}(\mathbf{z}, \mathbf{z}^i), \\ 190 \\ 191$$

192 where $\mathbf{z} = ((\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_{n+1}, \mathbf{y}_{n+1}))$ for calibration samples $((\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_n, \mathbf{y}_n))$ and test
 193 point $(\mathbf{x}_{n+1}, \mathbf{y}_{n+1})$ and \mathbf{z}^i arises from permuting the test point with the i^{th} calibration point.

194 For further details, we recommend the summary by Angelopoulos et al. (2024).

196

3.2 PDES AS OPERATOR MAPPINGS

198 Many dynamical systems in physics and engineering can be described by evolution equations of the
 199 form

$$200 \quad \frac{\partial u}{\partial t}(\mathbf{x}, t) = \mathcal{L}_{\mathbf{x}} u(\mathbf{x}, t), \\ 201$$

202 where $u : \Omega \times [0, \infty) \rightarrow \mathbb{R}$ is the state variable, $\mathbf{x} \in \Omega \subset \mathbb{R}^d$ denotes spatial coordinates, $t \geq 0$
 203 is the time, and $\mathcal{L}_{\mathbf{x}}$ is a (possibly nonlinear) differential operator acting on the spatial variable \mathbf{x} .
 204 We write $u_t := u(\cdot, t)$ for the spatial slice at time t . In this paper, we are interested in the Cauchy-
 205 type problem, where we consider boundary conditions on $\bar{\Omega}$ and initial conditions $u_0(\mathbf{x})$ from some
 206 Banach space of functions $(\mathcal{A}, \|\cdot\|_{\mathcal{A}})$ and are interested in a solution $u_t(\mathbf{x})$ in some Banach space
 207 of functions $(\mathcal{U}_t, \|\cdot\|_{\mathcal{U}_t})$. Typically, $u_t : \Omega \rightarrow \mathbb{R}$ and $u_t \in L^2(\Omega)$. We will further only consider
 208 well-posed problems, where we can define solution operators

$$209 \quad \mathcal{G}_t : \mathcal{A} \rightarrow \mathcal{U}_t, \quad \mathcal{G}_t(\mathbf{a})(\mathbf{x}) \mapsto u(\mathbf{x}, t) \\ 210$$

211 that uniquely map an initial condition to a solution function $u_t(\mathbf{x})$ and the map $t \mapsto \mathcal{G}_t$ is continuous
 212 in t . In the rest of the paper, we will assume that all functions come from the same space, so $\mathcal{A} = \mathcal{U}_t$
 213 for all $t \geq 0$, to simplify the notation, but the results apply more generally.

214 ¹The TV distance is originally defined on probability measures, and whenever we write $d_{\text{TV}}(\mathbf{x}, \mathbf{y})$ for
 215 random variables \mathbf{x} and \mathbf{y} , or $d_{\text{TV}}(\mathcal{P}_{\mathbf{x}}, \mathcal{P}_{\mathbf{y}})$ for probability distributions $\mathcal{P}_{\mathbf{x}}$ and $\mathcal{P}_{\mathbf{y}}$, we refer to the TV distance
 between their corresponding probability measures.

216 3.3 SURROGATE MODELS
217218 **Physics-Informed Neural Networks (PINNs).** PINNs (Raissi et al., 2019) approximate PDE
219 solutions by training a neural network to satisfy both observed data and the underlying PDE. The
220 loss function penalizes violations of the differential operator \mathcal{L} and boundary/initial conditions, so
221 that the neural network implicitly encodes the solution $u(\mathbf{x}, t)$. PINNs are flexible and require only
222 point-wise evaluations of the PDE residual, but they often struggle with stiff dynamics, sharp gradients,
223 or long time horizons.224 **Deep Operator Networks (DeepONets).** DeepONets (Lu et al., 2021) aim to directly learn non-
225 linear operators between function spaces. They decompose the problem into a *branch net*, which
226 encodes the input function (e.g., the initial condition), and a *trunk net*, which encodes the query point
227 (\mathbf{x}, t). The outputs are combined to approximate $u(\mathbf{x}, t) = \mathcal{G}_t(\mathbf{a})(\mathbf{x}, t)$. DeepONets provide a gen-
228 eral framework for operator learning and can handle diverse geometries and boundary conditions,
229 but require large and representative training data.230 **Neural Operators.** Neural operators Anandkumar et al. (2019); Li et al. (2021) generalize this
231 idea further by parameterizing mappings \mathcal{G} directly in function space, rather than through point-
232 wise regression. Unlike standard neural networks, which approximate finite-dimensional mappings,
233 neural operators approximate \mathcal{G} itself and can generalize across discretizations. In practice, functions
234 are observed on a finite set of points (grids or meshes), and the learned operator is evaluated on these
235 (or other) discretizations. Popular variants include the Fourier Neural Operator, which uses spectral
236 convolutions for global context, and the Graph Neural Operator, which extends to irregular meshes.237 **Other Surrogates.** Beyond these, there are also kernel-based approaches, reduced-order models,
238 and Gaussian process surrogates. However, in the machine learning literature, PINNs, DeepONets,
239 and neural operators have emerged as the three most prominent classes of PDE surrogates.240 4 WEIGHTED CONFORMAL PREDICTION FOR TIME-DEPENDENT PDE
241 SURROGATE MODELS

242 4.1 PROBLEM SETTING FOR CONFORMAL PREDICTION ON TIME-DEPENDENT PDES

243 To apply CP in the PDE setting, we start by specifying the underlying structure.

244 **From Initial Conditions to Solutions.** Assume we have an analytical form of the PDE, so that
245 we can generate our own data using numerical solvers. We first focus on the case where we want
246 to predict the solution at one fixed time point t for a given initial condition. To obtain our training
247 data $\mathcal{D}_{\text{train}}$, we would sample initial conditions $u_{0,i} \sim \mathcal{P}_0$, $i = 1, \dots, N_{\text{train}}$, from a distribution
248 on \mathcal{U} , and obtain the corresponding solution at time t by numerically solving the PDE. This defines
249 a pushforward measure²

250
$$\mathcal{P}_t := (S_t)_\# \mathcal{P}_0,$$

251 where $S_t : \mathcal{U} \rightarrow \mathcal{U}$ is the PDE solution operator mapping initial conditions u_0 to solutions u_t . Our
252 training dataset then consists of

253
$$\mathcal{D}_{\text{train}} = \{(u_{0,i}, u_{t,i})\}_{i=1}^{N_{\text{train}}}, \quad u_{t,i} \sim \mathcal{P}_t.$$

254 If we now consider consecutive time points, our distribution changes over time:

255
$$\mathcal{P}_0 \xrightarrow{S_\delta} \mathcal{P}_\delta \xrightarrow{S_\delta} \mathcal{P}_{2\delta} \xrightarrow{S_\delta} \dots$$

256 Thus, we obtain a sequence of probability distributions $\{\mathcal{P}_t\}_{t \geq 0}$ on the same function space, evolv-
257 ing under the PDE dynamics.258 **Implication for Conformal Prediction.** Calibration and test data drawn from different \mathcal{P}_t are
259 therefore **not exchangeable**: although they live in the same function space, their distributions shift
260 with time.261 ²We slightly abuse notation here by writing the pushforward in terms of the distribution instead of the
262 measure corresponding to the distribution.

270 4.2 DISTRIBUTION SHIFTS IN FUNCTION SPACES
271

272 Having specified the problem setup, we now investigate if we can calculate the TV distance between
273 the laws of a PDE solution at different time points. If the TV distance of the laws of time points t
274 and $t + \delta$ were moderate, we could recover CP coverage guarantees for the $t + \delta$ prediction using
275 the approach from Barber et al. (2023).

276 We will start by analyzing the problem in the function-space setting, as is often employed in the
277 neural operator literature and related CP works (Harris & Liu, 2025; Gray et al., 2025; Mollaali
278 et al., 2024). We will show that even for a simple PDE, like the heat equation with Gaussian initial
279 distribution, the TV distance between the solution-distributions $\mathcal{P}_t, \mathcal{P}_{t+\delta}$ at two time points $t, t + \delta$
280 is always maximal,

$$281 d_{\text{TV}}(\mathcal{P}_t, \mathcal{P}_{t+\delta}) = 1, \quad \text{for all } t \geq 0, \delta > 0.$$

282 This is representative of a broader phenomenon that “[...] measures in infinite-dimensional spaces
283 have a strong tendency of being mutually singular.” Hairer (2023). **As a direct consequence, regular**
284 **CP—and any method relying on equality or even approximate similarity between calibration and test**
285 **distributions—becomes inapplicable. Regaining guarantees would require stronger implicit biases,**
286 **but this lies beyond the scope of this paper.**

287 Finally, note that, while this issue complicates theoretical considerations in the neural operator literature,
288 it is not necessarily problematic for practical CP on surrogate models. In practice, we always
289 work with finite-dimensional discretizations, which mitigate this effect, as will be discussed in section
290 4.3.

291 **Theorem 4.1.** *Consider the one-dimensional heat equation on the domain $\Omega = (0, 1)$ with Dirichlet
292 boundary conditions*

$$293 \frac{\partial u}{\partial t}(x, t) = \frac{\partial^2 u}{\partial x^2}(x, t), \quad x \in (0, 1), t \geq 0, \\ 294 u(0, t) = u(1, t) = 0, \quad t \geq 0, \\ 295 u(x, 0) = u_0(x), \quad x \in (0, 1),$$

296 where $u : \bar{\Omega} \times [0, \infty) \rightarrow \mathbb{R}$ denotes the temperature at location x and time t . Suppose the initial
297 condition is sampled from a Gaussian distribution

$$298 \mathcal{P}_0 \sim \mathcal{N}(\mathbf{0}, (\mathbf{I} - \mathbf{\Lambda})^{-1}),$$

299 where $\mathbf{\Lambda}$ is the Laplace operator on Ω with Dirichlet boundary conditions. Then, for any $t \geq 0$,
300 $\delta > 0$, the TV distance between the measures \mathcal{P}_t and $\mathcal{P}_{t+\delta}$ of the solution $u(\cdot, t)$ and $u(\cdot, t + \delta)$ is
301 maximal, i.e.

$$302 d_{\text{TV}}(\mathcal{P}_t, \mathcal{P}_{t+\delta}) = 1.$$

303 The proof is provided in appendix A.2.

304 We will now discuss how, despite the issue above, coverage guarantees can be recovered for time-
305 dependent PDE surrogate models in practice.

306 4.3 RECOVERING COVERAGE GUARANTEES
307

308 The following theorem provides the exact distribution of the solution u_t on a discretized space, using
309 the *method of lines*. We provide an intuitive example in appendix A.3.

310 **Theorem 4.2.** *Let $\Omega \subset \mathbb{R}^d$ be a bounded domain, and let*

$$311 \mathcal{M} := \{x_1, \dots, x_n\} \subset \Omega$$

312 denote a discretization of Ω . Consider the finite-difference scheme in space, with $\mathbf{A} \in \mathbb{R}^{n \times n}$ ap-
313 proximating the solution of

$$314 \frac{\partial u}{\partial t}(\mathbf{x}, t) = \mathcal{L}_{\mathbf{x}} u(\mathbf{x}, t), \quad \mathbf{x} \in \Omega, t \geq 0,$$

315 with linear boundary conditions on $\partial\Omega$, where $\mathcal{L}_{\mathbf{x}}$ is a linear spatial differential operator. This
316 yields the discretized dynamics

324

$$\frac{d\mathbf{u}(t)}{dt} = \mathbf{A}\mathbf{u}(t) + \mathbf{r}(t), \quad \mathbf{u}(t), \mathbf{r}(t) \in \mathbb{R}^n.$$

327

328 Suppose the initial condition satisfies $\mathbf{u}(0) \sim \mathcal{N}(\mu_0, \Sigma_0)$. Then, for $t \geq 0$ and $\delta > 0$, the law \mathcal{P}_t of
 329 $\mathbf{u}(t)$ is Gaussian with mean

330

$$\mu_t = \exp(t\mathbf{A})\mu_0 + \int_0^t \exp((t-s)\mathbf{A})\mathbf{r}(s)ds$$

333

334 and covariance

335

$$\Sigma_t = \exp(t\mathbf{A})\Sigma_0 \exp(t\mathbf{A}^T).$$

337

338 *Proof.* As we discretized only in space, not in time, the finite difference scheme yields a linear
 339 system of ODEs

340

$$\frac{d\mathbf{u}(t)}{dt} = \mathbf{A}\mathbf{u}(t) + \mathbf{r}(t).$$

341

342 As \mathbf{A} is independent of t and $\mathbf{r}(t)$ is the deterministic source term, the solution of the system of
 343 ODEs is given by

344

$$\mathbf{u}(t) = \exp(t\mathbf{A})\mathbf{u}(0) + \int_0^t \exp((t-s)\mathbf{A})\mathbf{r}(s)ds.$$

345

346 Note that we assumed $\mathbf{u}(0)$ is Gaussian, i.e.,

347

$$\mathbf{u}_0 \sim \mathcal{N}(\mu_0, \Sigma_0), \quad \mu_0 \in \mathbb{R}^n, \Sigma_0 \in \mathbb{R}^{n \times n},$$

348

349 and $\exp(t\mathbf{A})$ is just a matrix, so $\mathbf{u}(t)$ is also Gaussian with mean $\mu_t = \exp(t\mathbf{A})\mu_0 + \int_0^t \exp((t-s)\mathbf{A})\mathbf{r}(s)ds$ and covariance $\Sigma_t = \exp(t\mathbf{A})\Sigma_0 \exp(t\mathbf{A}^T)$.
 350

351

□

352

353 *Remark 4.3.* This result can be generalized to other initial distributions. The location-scale family
 354 of distributions, for example, is closed under affine transformations leading to similar results.
 355 The location-scale family includes, among others, the Gaussian, Cauchy, Laplace, and logistic distributions.
 356 Note, however, that the Gaussian assumption we made is the most common in recent literature
 357 (Li et al., 2021; Santos et al., 2023; Gopakumar et al., 2024; Zhou & Barati Farimani,
 358 2025; Gopakumar et al., 2025). Also, from a physical viewpoint, a Gaussian random field aligns
 359 well with the laws of nature in the sense that the aggregate effect of many small independent perturbations,
 360 forming the initial condition, is approximately Gaussian by the central limit theorem.
 361 [We added additional experiments, sampling from different location-scale initial distributions in the appendix A.8.](#)

362

363 *Remark 4.4.* Theorem 4.2 also allows us to derive an upper bound on the TV distance of the laws of
 364 u_t and $u_{t+\delta}$. While we will not make use of this result in our method, we provide the theorem and
 365 proof in appendix A.4.

366

4.4 LIKELIHOOD-WEIGHTED CONFORMAL PREDICTION

367

368 Theorem 4.2 shows that under a discretized linear PDE with Gaussian initial conditions, the solution
 369 at time t is Gaussian with mean μ_t and covariance Σ_t as stated in the theorem. Consequently, both
 370 calibration and test distributions (corresponding to time points t and $t + \delta$ for one or more $\delta > 0$)
 371 are Gaussian and their density ratio is available in closed form. This enables a likelihood-weighted
 372 conformal predictor:

373

$$w_{i,\delta} \propto \frac{\mathcal{N}(\mathbf{u}_i; \mu_{t+\delta}, \Sigma_{t+\delta})}{\mathcal{N}(\mathbf{u}_i; \mu_t, \Sigma_t)}, \quad (1)$$

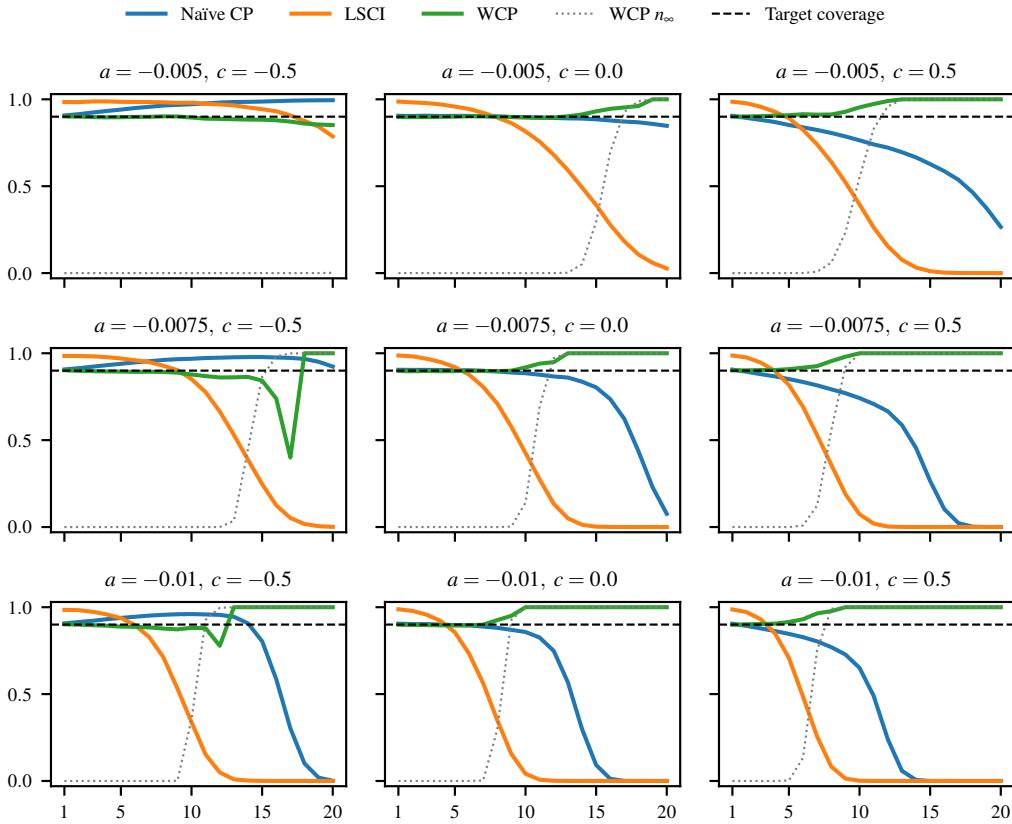
375

376 for all \mathbf{u}_i belonging to the calibration set together with the target test point. Normalizing these
 377 weights and applying split CP with the weighted quantile yields conformal bands with formal cov-
 erage guarantees.

378
 379 *Remark 4.5.* Within this CP framework applied to the discretized setting, we provide asymptotic—
 380 and in some cases even non-asymptotic—guarantees for the PDE solution $u(x, t)$ in the original
 381 space. The nature of the bounds depends on both the PDE and the discretization scheme, but the
 382 key idea is that the bands on the discretized solution can be transferred to the original solution by
 383 leveraging numerical error guarantees of the scheme.

384 **5 EXPERIMENTS**
 385

386 **[NEW FIGURE]**
 387



416 Figure 3: Mean coverages for varying a, c and fixed $b = -0.5$ across increasing prediction horizon.
 417 We omit coverages when infinite conformal bands were reported (coverage of 1 would hold trivially).
 418 The 90% coverage target and, for WCP, the fraction of samples with infinite bandwidth n_∞ are
 419 marked as black dashed line and gray dotted line, respectively.

420
 421 **Experimental Setup** We demonstrate our method on both synthetic and real-world data. For the
 422 synthetic case, we choose the general second order PDE framework

$$424 \quad u_t(x, t) + au_{xx}(x, t) + bu_x(x, t) + cu(x, t) = 0, \quad x \in (0, 1), t \geq 0, \\ 425 \quad u(0, t) = u(1, t) = 0, \quad t \geq 0, \\ 426 \quad u(x, 0) = u_0(x), \quad x \in (0, 1).$$

427 where $a, b, c \in \mathbb{R}$ are tuneable parameters. To test WCP, we target setups where the PDE becomes
 428 more unstable over time ($a < 0$)—otherwise, CP methods that rely on observed residuals trivially
 429 cover at future time steps. Hence, we consider $a = \{-0.005, -0.0075, -0.01\}$ and for the remaining
 430 parameters we choose $b, c \in \{-0.5, 0, 0.5\}$. As a base model, we train a geometry-informed
 431 neural operator (Li et al., 2023) and calibrate on the residuals with the respective CP method (note

432 Table 1: Mean coverages and bandwidths over 5000 sampled initial conditions for varying a and
 433 fixed $b = -0.5$, $c = -0.5$. For WCP, we also report the fraction of samples where infinite bands
 434 were reported (n_∞) to maintain coverage guarantees. The gray font is chosen for better readability.
 435

			Timestep				
			1	5	10	15	20
$a = -0.005$	Naïve CP	Coverage	0.91	0.94	0.97	0.99	0.99
		Bandwidth	0.03	0.03	0.03	0.03	0.03
	LSCI	Coverage	0.98	0.99	0.98	0.94	0.79
		Bandwidth	0.02	0.02	0.02	0.02	0.02
	WCP (Ours)	Coverage	0.9	0.9	0.9	0.88	0.85
		Bandwidth	0.03	0.03	0.03	0.02	0.02
		n_∞	0.0%	0.0%	0.0%	0.0%	0.2%
$a = -0.0075$	Naïve CP	Coverage	0.91	0.94	0.97	0.98	0.92
		Bandwidth	0.03	0.03	0.03	0.03	0.03
	LSCI	Coverage	0.98	0.97	0.85	0.25	0.0
		Bandwidth	0.02	0.02	0.02	0.02	0.02
	WCP (Ours)	Coverage	0.9	0.89	0.88	0.84	1.0
		Bandwidth	0.03	0.03	0.03	0.03	∞
		n_∞	0.0%	0.0%	0.0%	86.4%	100%
$a = -0.01$	Naïve CP	Coverage	0.91	0.94	0.96	0.8	0.0
		Bandwidth	0.03	0.03	0.03	0.03	0.03
	LSCI	Coverage	0.98	0.94	0.34	0.0	0.0
		Bandwidth	0.02	0.02	0.02	0.02	0.02
	WCP (Ours)	Coverage	0.9	0.89	0.88	1.0	1.0
		Bandwidth	0.03	0.03	0.03	∞	∞
		n_∞	0.0%	0.0%	35.4%	100%	100%

462 that the choice of surrogate model is not important for downstream analysis). The task of the base
 463 model is to predict the solution u_t at 20 time steps in the future. The task of the CP methods is to
 464 report conformal bands with 90% coverage. For each PDE, we sample 5000 trajectories to train the
 465 base model, 500 for validation, and 5000 for calibration and testing each. We adjusted the time steps
 466 and other parameters individually with more details in appendix A.5.

467 **Baselines** We define two baselines for our experiments. The first is a naïve implementation with
 468 no consideration of exchangeability (naïve CP). Specifically, we implemented Diquigiovanni et al.
 469 (2022), who define the score as the maximum absolute error over space and use the regular split CP
 470 algorithm. Since exchangeability does not hold in this setup, the conformal bands of naïve CP have
 471 **no formal guarantees**.

472 Secondly, we use the LSCI method (Harris & Liu, 2025) ($\lambda = 3$, projection dimension: 20, number
 473 of CP band samples: 5000). We choose a large number of band samples to push LSCI to over-
 474 coverage, so undercoverage can be evaluated in a fair manner. Note that because their guarantees
 475 only hold under the local exchangeability assumption which is not verifiable (see appendix A.1), the
 476 LSCI CP bands also have **no formal guarantees** in our experiments.

477 Our weighted conformal prediction (WCP) method is based on a weighted version of Diquigiovanni
 478 et al. (2022). Specifically, knowing that our solution is Gaussian at every time point, we weigh our
 479 score according to equation (1).

480 **Evaluation** For each method and each PDE, we report the mean coverage and bandwidth of the
 481 5000 test set samples. We consider a sample covered if all of points of the function are within the
 482 conformal bands. In cases where the distributional dissimilarity of u_t and $u_{t+\delta}$ is too large, our
 483 WCP method predicts infinite bands. If this is the case, we exclude the sample and only predict
 484 coverage of the other samples. We report the fraction of excluded samples n_∞ in our results.

486 Note that reporting trivial bands is usually a more valuable result than delivering bands with un-
 487 dercoverage, especially in safety-critical tasks. The key strength of CP is its coverage guarantees
 488 and our WCP detects when it cannot predict meaningful bands and refrains from violating the target
 489 coverage.
 490

491

492

493 **Results** We report results for varying a, c with $b = -0.5$ in Figure 3, and provide the correspond-
 494 ing plots for the remaining b -values in appendix A.7. For $b = c = -0.5$, the numerical results are
 495 listed in Table 1, while results for the other c -values are given in appendix A.7, together with a visu-
 496 alization of CP bands. Overall, in most configurations, naïve CP and LSCI fail to meet the coverage
 497 target—earlier and more severely as the PDE becomes noisier (i.e., for smaller a)—while WCP con-
 498 sistently meets its coverage guarantees. When n_∞ approaches roughly 90%, WCP shows a slight
 499 drop in empirical coverage. This behavior is expected, as we only report coverage of non-trivial
 500 bands: with very few samples remaining, the empirical coverage is subject to higher stochastic
 501 noise. In practice, this can be addressed by using the bands only for sufficiently large remaining
 502 sample size or by considering the overall coverage including the trivial bands.
 503

504 As discussed above, our method reports infinite bands for increasing distribution shift. Although
 505 this sacrifices meaningful bands, it ensures fully reliable coverage guarantees. Lastly, we observed
 506 that WCP and naïve CP are significantly faster than LSCI: When running LSCI on a MacBook Pro
 507 M4 Pro with 24GB RAM, sampling the conformal bands for 5000 test samples takes approximately
 508 40 minutes. The WCP and the naïve method take only seconds. Overall, WCP is the only method
 509 providing **formal guarantees**, and we can see empirically that this is a clear advantage as soon as
 510 our system exhibits significant dynamics.
 511

512 **Real-World Example** To demonstrate the applicability of our method in real-world scenarios,
 513 we use the dataset of Wei et al. (2023). They provide a small 2D-dataset of pulsed-thermography
 514 measurements, where objects are heated and then cooled while surface temperatures are recorded
 515 to detect subsurface defects. We use only the cooldown phase, as it approximately follows the heat
 516 equation. We provide more details on our implementations and the results in appendix A.6. Our
 517 method achieves target coverage over all tested time steps.
 518

519

6 DISCUSSION

520

521 Conformal prediction for time-dependent physical phenomena is often constrained by non-
 522 exchangeable data. In this work, we investigated whether coverage guarantees can be maintained
 523 beyond the exchangeability assumption. Our results show that this depends strongly on the setup.
 524 On function spaces, measures are typically mutually singular, making coverage guarantees unattain-
 525 able. On discretized domains, however, we derived how weighted CP can be applied to linear PDEs
 526 to obtain coverage guarantees. We empirically validated that weighted CP is the only method that
 527 reliably achieves the target coverage compared to baselines.
 528

529

530

531 These findings connect back to our starting point: non-stationarity in time-dependent PDEs breaks
 532 classical CP, but weighted CP offers a principled alternative. We established coverage for the class of
 533 linear PDEs. Although this class covers many practical problems, extending the analysis to nonlinear
 534 PDEs is a natural next step and would further broaden the applicability of conformal prediction in
 535 scientific machine learning.
 536

537

538

REPRODUCIBILITY STATEMENT

539

We provided the code for the data generation, model training, fitting of conformal bands, and instructions on how to run it as supplementary material to the reviewers. With that, all figures and results can be reproduced independently. For the final version, we will set up a public GitHub repository. The proof for theorem 4.2 can be found in the main text, and the proof for theorem 4.1 can be found in appendix A.2.

540 REFERENCES
541

542 Anima Anandkumar, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Nikola Kovachki, Zongyi Li,
543 Burigede Liu, and Andrew Stuart. Neural Operator: Graph Kernel Network for Partial Differential
544 Equations. In *ICLR 2020 Workshop on Integration of Deep Neural Models and Differential
545 Equations*, 2019. URL <https://openreview.net/forum?id=fg2ZFmXF03>.

546 Anastasios N. Angelopoulos, Emmanuel J. Candes, and Ryan J. Tibshirani. Conformal PID Control
547 for Time Series Prediction, August 2023. URL <http://arxiv.org/abs/2307.16895> [cs].
548 arXiv:2307.16895 [cs].

549 Anastasios Nikolas Angelopoulos, Rina Foygel Barber, and Stephen Bates. Theoretical Foundations
550 of Conformal Prediction. 2024. URL <https://api.semanticscholar.org/CorpusID:274131633>.
551

552 Rina Foygel Barber, Emmanuel J. Candes, Aaditya Ramdas, and Ryan J. Tibshirani. Conformal
553 prediction beyond exchangeability, March 2023. URL <http://arxiv.org/abs/2202.13415> [stat].
554

555 Giuseppe Da Prato and Jerzy Zabczyk. *Stochastic Equations in Infinite Dimensions*. Cambridge
556 University Press, 1 edition, December 1992. ISBN 978-0-521-38529-9 978-0-521-05980-0 978-
557 0-511-66622-3. doi: 10.1017/CBO9780511666223. URL <https://www.cambridge.org/core/product/identifier/9780511666223/type/book>.
558

559 Jacopo Diquigiovanni, Matteo Fontana, and Simone Vantini. Conformal prediction bands for
560 multivariate functional data. *Journal of Multivariate Analysis*, 189:104879, May 2022. ISSN
561 0047259X. doi: 10.1016/j.jmva.2021.104879. URL <https://linkinghub.elsevier.com/retrieve/pii/S0047259X21001573>.
562

563 Isaac Gibbs and Emmanuel Candès. Adaptive Conformal Inference Under Distribution Shift, De-
564 cember 2021. URL <http://arxiv.org/abs/2106.00170>. arXiv:2106.00170 [stat].
565

566 Vignesh Gopakumar, Ander Gray, Joel Oskarsson, Lorenzo Zanisi, Stanislas Pamela, Daniel Giles,
567 Matt J. Kusner, and Marc Peter Deisenroth. Uncertainty Quantification of Surrogate Models using
568 Conformal Prediction. 2024. URL <https://api.semanticscholar.org/CorpusID:271903374>.
569

570 Vignesh Gopakumar, Ander Gray, Lorenzo Zanisi, Timothy Nunn, Daniel Giles, Matt J. Kusner,
571 Stanislas Pamela, and Marc Peter Deisenroth. Calibrated Physics-Informed Uncertainty Quantifi-
572 cation, 2025. URL <https://arxiv.org/abs/2502.04406>. Version Number: 2.
573

574 Ander Gray, Vignesh Gopakumar, Sylvain Rousseau, and Sébastien Destercke. Guaranteed Predic-
575 tion Sets for Functional Surrogate Models. In *Conference on Uncertainty in Artificial Intelligence*,
576 2025. URL <https://api.semanticscholar.org/CorpusID:275993804>.
577

578 Martin Hairer. An Introduction to Stochastic PDEs, July 2023. URL <http://arxiv.org/abs/0907.4178>. arXiv:0907.4178 [math].
579

580 Trevor Harris and Yan Liu. Locally Adaptive Conformal Inference for Operator Models, July 2025.
581 URL <http://arxiv.org/abs/2507.20975>. arXiv:2507.20975 [stat].
582

583 Zongyi Li, Nikola Borislavov Kovachki, Kamyar Azizzadenesheli, Burigede liu, Kaushik Bhat-
584 charya, Andrew Stuart, and Anima Anandkumar. Fourier Neural Operator for Parametric Par-
585 tial Differential Equations. In *International Conference on Learning Representations*, 2021. URL
586 <https://openreview.net/forum?id=c8P9NQVtmnO>.
587

588 Zongyi Li, Nikola Kovachki, Chris Choy, Boyi Li, Jean Kossaifi, Shourya Otta, Moham-
589 mad Amin Nabian, Maximilian Stadler, Christian Hundt, Kamyar Azizzadenesheli, and An-
590 imashree Anandkumar. Geometry-Informed Neural Operator for Large-Scale 3D PDEs. In
591 A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in
592 Neural Information Processing Systems*, volume 36, pp. 35836–35854. Curran Associates, Inc.,
593 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/70518ea42831f02afc3a2828993935ad-Paper-Conference.pdf.
594

594 Lu Lu, Pengzhan Jin, and George Em Karniadakis. DeepONet: Learning nonlinear operators
 595 for identifying differential equations based on the universal approximation theorem of operators.
 596 *Nature Machine Intelligence*, 3(3):218–229, March 2021. ISSN 2522-5839. doi: 10.1038/s42256-021-00302-5. URL <http://arxiv.org/abs/1910.03193>. arXiv:1910.03193
 597 [cs].
 598

599 Amirhossein Mollaali, Gabriel Zufferey, Gonzalo E. Constante-Flores, Christian Moya, Can Li,
 600 Guang Lin, and Meng Yue. Conformalized Prediction of Post-Fault Voltage Trajectories Using
 601 Pre-trained and Finetuned Attention-Driven Neural Operators. *Neural networks : the official
 602 journal of the International Neural Network Society*, 192:107809, 2024. URL <https://api.semanticscholar.org/CorpusID:273707610>.
 603

604 Christian Moya, Amirhossein Mollaali, Zecheng Zhang, Lu Lu, and Guang Lin. Conformalized-
 605 DeepONet: A distribution-free framework for uncertainty quantification in deep operator net-
 606 works. *Physica D: Nonlinear Phenomena*, 471:134418, January 2025. ISSN 01672789. doi: 10.
 607 1016/j.physd.2024.134418. URL <https://linkinghub.elsevier.com/retrieve/pii/S0167278924003683>.
 608

609 Roberto I. Oliveira, Paulo Orenstein, Thiago Ramos, and João Vitor Romano. Split Conformal
 610 Prediction and Non-Exchangeable Data. *J. Mach. Learn. Res.*, 25:225:1–225:38, 2022. URL
 611 <https://api.semanticscholar.org/CorpusID:247794086>.
 612

613 M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning
 614 framework for solving forward and inverse problems involving nonlinear partial differential equa-
 615 tions. *Journal of Computational Physics*, 378:686–707, February 2019. ISSN 00219991. doi:
 616 10.1016/j.jcp.2018.10.045. URL <https://linkinghub.elsevier.com/retrieve/pii/S0021999118307125>.
 617

618 Fabricio Dos Santos, Tara Akhound-Sadegh, and Siamak Ravanbakhsh. Physics-Informed Trans-
 619 former Networks. In *The Symbiosis of Deep Learning and Differential Equations III*, 2023. URL
 620 <https://openreview.net/forum?id=zu80h9YryU>.
 621

622 Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random
 623 World*. Springer International Publishing, Cham, 2022. ISBN 978-3-031-06648-1 978-3-031-
 624 06649-8. doi: 10.1007/978-3-031-06649-8. URL <https://link.springer.com/10.1007/978-3-031-06649-8>.
 625

626 Ziang Wei, Ahmad Osman, Bernd Valeske, and Xavier Maldague. Pulsed thermography dataset
 627 for training deep learning models. *Applied Sciences*, 13(5), 2023. ISSN 2076-3417. doi: 10.
 628 3390/app13052901. URL <https://www.mdpi.com/2076-3417/13/5/2901>. Number:
 629 2901.
 630

631 Chen Xu and Yao Xie. Sequential Predictive Conformal Inference for Time Series, May 2023. URL
 632 <http://arxiv.org/abs/2212.03463>. arXiv:2212.03463 [stat].
 633

634 Anthony Zhou and Amir Barati Farimani. Predicting change, not states: An alternate framework for
 635 neural PDE surrogates. *Computer Methods in Applied Mechanics and Engineering*, 441:117990,
 636 June 2025. ISSN 00457825. doi: 10.1016/j.cma.2025.117990. URL <https://linkinghub.elsevier.com/retrieve/pii/S0045782525002622>.
 637

638

639 A APPENDIX

641 A.1 VALIDATING LOCAL EXCHANGEABILITY

643 To have formal guarantees with the LSCI approach from Harris & Liu (2025), it is required that the
 644 model’s residuals r_t are locally exchangeable. Considering this for the most simple setup of two
 645 time points t and $t + \delta$ for some $\delta > 0$, this means that it must hold that
 646

$$647 d_{\text{TV}}\left(\begin{pmatrix} r_t \\ r_{t+\delta} \end{pmatrix}, \begin{pmatrix} r_{t+\delta} \\ r_t \end{pmatrix}\right) \leq d(t, t + \delta)$$

for a pre-metric d on the time domain. Note however, that since we do not have access to the laws of the random vectors above, we cannot reason about their TV distance. Even though in theorem 4.2 we derive the laws of the solutions u_t , it is not clear how to reason about the above left hand side without further assumptions (like independence of the residuals over time—which is not plausible as the solution has a clear time-dependence in all non-stationary PDEs).

A.2 PROOF OF THEOREM 4.1

Theorem A.1. *Consider the one-dimensional heat equation on the domain $\Omega = (0, 1)$ with Dirichlet boundary conditions*

$$\begin{aligned} \frac{\partial u}{\partial t}(x, t) &= \frac{\partial^2 u}{\partial x^2}(x, t), \quad x \in (0, 1), t \geq 0, \\ u(0, t) &= u(1, t) = 0, \quad t \geq 0, \\ u(x, 0) &= u_0(x), \quad x \in (0, 1), \end{aligned}$$

where $u : \bar{\Omega} \times [0, \infty) \rightarrow \mathbb{R}$ denotes the temperature at location x and time t . Suppose we sample the initial condition from a Gaussian distribution

$$\mathcal{P}_0 \sim \mathcal{N}(\mathbf{0}, (\mathbf{I} - \mathbf{\Lambda})^{-1}),$$

where $\mathbf{\Lambda}$ is the Laplace operator on Ω with Dirichlet boundary conditions. Then, for any $t \geq 0$, $\delta > 0$, the TV distance between the measures \mathcal{P}_t and $\mathcal{P}_{t+\delta}$ of the solution $u(\cdot, t)$ and $u(\cdot, t + \delta)$ is maximal, i.e.

$$d_{\text{TV}}(\mathcal{P}_t, \mathcal{P}_{t+\delta}) = 1.$$

Proof. Notice that our whole setup is on a Hilbert Space. We begin by showing that the $(\mathbf{I} - \mathbf{\Lambda})^{-1}$ is a well-defined covariance operator on $L^2(\Omega)$. For this, according to Hairer (2023)[Proposition 4.17], the operator must be positive, symmetric and trace class. As both \mathbf{I} and $-\mathbf{\Lambda}$ are positive and symmetric, so is their sum. As $-\mathbf{\Lambda}$ is positive, it follows that $(\mathbf{I} - \mathbf{\Lambda})$ is invertible, due to strictly positive eigenvalues. Further, the eigenvalues of $-\mathbf{\Lambda}$ with Dirichlet boundary conditions are given by $\lambda_n = (n\pi)^2$, with corresponding eigenfunctions $e_n(x) = \sqrt{2} \sin(n\pi x)$, which form an orthonormal basis of $L^2(\Omega)$. Therefore, the eigenvalues of $(\mathbf{I} - \mathbf{\Lambda})^{-1}$ are given by $\mu_n = 1/(1 + (n\pi)^2)$ and as

$$\sum_{n=1}^{\infty} \mu_n < \infty,$$

we conclude that $(\mathbf{I} - \mathbf{\Lambda})^{-1}$ is a trace class operator and thus defines a Gaussian measure on $L^2(\Omega)$.

Now our proof will be based on the Feldman-Hájek theorem Da Prato & Zabczyk (1992)[Theorem 2.23], which gives a characterization of when two Gaussian measures on a Hilbert space are either equivalent or mutually singular. We will briefly state the whole chain of reasoning, and then provide the necessary details.

We will show that our measure at all times is Gaussian. By the Feldman-Hájek theorem, two Gaussian measures $\mathcal{N}(\mathbf{m}_1, \mathbf{C}_1)$ and $\mathcal{N}(\mathbf{m}_2, \mathbf{C}_2)$ on a Hilbert space are either equivalent or mutually singular. A necessary condition for equivalence is that the Cameron-Martin spaces, as given by $\mathbf{C}^{1/2}$, of the two measures are equal as sets Da Prato & Zabczyk (1992)[Theorem 2.23]. Thus, if the ranges of the covariance operators $\mathbf{C}_1^{1/2}$ and $\mathbf{C}_2^{1/2}$ are not equal, then the measures are mutually singular and their TV distance is 1.

Calculating the Covariance Operators Starting with a measure μ_0 of the initial distribution, the heat equation induces a semigroup $S(t) = \exp(t\mathbf{\Lambda})$, which maps the initial condition u_0 to the solution at time t , i.e. $u(\cdot, t) = S(t)u_0$. Therefore, the measure μ_t of $u(\cdot, t)$ is induced by the pushforward measure μ_0 under $S(t)$, i.e. $\mu_t = S(t)\#\mu_0$. As $S(t)$ is linear, μ_t is also a Gaussian measure with mean $\mathbf{0}$ and covariance operator

$$\mathbf{C}_t = S(t)(\mathbf{I} - \mathbf{\Lambda})^{-1}S(t)^* = \exp(t\mathbf{\Lambda})(\mathbf{I} - \mathbf{\Lambda})^{-1}\exp(t\mathbf{\Lambda}),$$

where $S(t)^*$ denotes the adjoint of $S(t)$ (Hairer (2023) Chap. 4.3). As we have seen above, the eigenvalues of $(\mathbf{I} - \mathbf{\Lambda})^{-1}$ are given by $\mu_n = 1/(1 + (n\pi)^2)$, with corresponding eigenfunctions

702 $e_n(x)$. Further, the eigenfunctions of Λ are also given by $e_n(x)$, with corresponding eigenvalues
 703 $\lambda_n = -(n\pi)^2$. Lastly, by functional calculus, the eigenfunctions of $\exp(t\Lambda)$ are also given by $e_n(x)$
 704 with corresponding eigenvalues $\nu_n = \exp(-t(n\pi)^2)$. With this, we can compute

$$705 \quad \mathbf{C}_t \mathbf{e}_n = \exp(t\Lambda)(\mathbf{I} - \Lambda)^{-1} \exp(t\Lambda) \mathbf{e}_n \\ 706 \quad = (\exp(-t(n\pi)^2)) \left(\frac{1}{1 + (n\pi)^2} \right) (\exp(-t(n\pi)^2)) \mathbf{e}_n = \frac{\exp(-2t(n\pi)^2)}{1 + (n\pi)^2} \mathbf{e}_n.$$

709 Thus, the eigenvalues of \mathbf{C}_t are given by $\lambda_n(t) = \nu_n^2 \mu_n = \exp(-2t(n\pi)^2)/(1 + (n\pi)^2)$, with
 710 corresponding eigenfunctions $e_n(x)$.
 711

712 **Calculating the Cameron-Martin Spaces** The functions $e_n(x)$ form an orthonormal basis of
 713 $L^2(\Omega)$, so we can express every element $f \in L^2(\Omega)$ as

$$714 \quad f = \sum_{n=1}^{\infty} c_n e_n, \quad \sum_{n=1}^{\infty} c_n^2 < \infty.$$

717 The Cameron-Martin space H_t of \mathcal{P}_t is given by the range of $\mathbf{C}_t^{1/2}$, which is given by
 718

$$719 \quad \text{Ran}(\mathbf{C}_t^{1/2}) = \left\{ \mathbf{C}_t^{1/2} f \mid f \in L^2(\Omega) \right\} = \left\{ \sum_{n=1}^{\infty} \sqrt{\lambda_n(t)} c_n e_n \mid f \in L^2(\Omega) \right\}.$$

721 Therefore, $g \in H_t$ if and only if g can be expressed as
 722

$$723 \quad g = \sum_{n=1}^{\infty} d_n e_n, \quad \sum_{n=1}^{\infty} \frac{d_n^2}{\lambda_n(t)} < \infty.$$

725 Inserting the expression for $\lambda_n(t)$, we see that $g \in H_t$ if and only if
 726

$$727 \quad \sum_{n=1}^{\infty} d_n^2 \frac{1 + (n\pi)^2}{\exp(-2t(n\pi)^2)} < \infty.$$

730 **Showing That the Cameron-Martin Spaces Are Not Equal** Now it is easy to see that for any
 731 $t \geq 0, \delta > 0$, the Cameron-Martin spaces H_t and $H_{t+\delta}$ are not equal. For example, the function
 732

$$733 \quad h(x) = \sum_{n=1}^{\infty} \exp(-(t + \delta)(n\pi)^2) e_n(x)$$

735 is an element of H_t , as

$$736 \quad \sum_{n=1}^{\infty} (\exp(-(t + \delta)(n\pi)^2))^2 \frac{1 + (n\pi)^2}{\exp(-2t(n\pi)^2)} = \sum_{n=1}^{\infty} (1 + (n\pi)^2) \exp(-2\delta(n\pi)^2) < \infty,$$

739 but it is not an element of $H_{t+\delta}$, as

$$740 \quad \sum_{n=1}^{\infty} (\exp(-(t + \delta)(n\pi)^2))^2 \frac{1 + (n\pi)^2}{\exp(-2(t + \delta)(n\pi)^2)} = \sum_{n=1}^{\infty} (1 + (n\pi)^2) = \infty.$$

743 Therefore, by the Feldman-Hájek theorem, the measures \mathcal{P}_t and $\mathcal{P}_{t+\delta}$ are mutually singular, and
 744 their TV distance is 1. \square
 745

746 A.3 ILLUSTRATION OF THE METHOD OF LINES

748 Consider the one-dimensional heat equation on the domain $\Omega = (0, 1)$ with Dirichlet boundary
 749 conditions

$$750 \quad \frac{\partial u}{\partial t}(x, t) = \frac{\partial^2 u}{\partial x^2}(x, t), \quad x \in (0, 1), t \geq 0, \\ 751 \quad u(0, t) = u(1, t) = 0, \quad t \geq 0, \\ 752 \quad u(x, 0) = u_0(x), \quad x \in (0, 1),$$

755 where $u : \bar{\Omega} \times [0, \infty) \rightarrow \mathbb{R}$ denotes the temperature field. We will numerically solve this PDE using
 the *method of lines*.

756 **Method of Lines** We discretize the spatial domain with a uniform grid $\mathcal{M} = \{x_1, \dots, x_n\} \subset \Omega$
 757 with $x_i = \frac{i}{n+1}, i \in \{1, \dots, n\}$, while leaving the time domain continuous. We can approximate the
 758 second derivative in space with the finite difference scheme
 759

$$760 \quad \frac{\partial^2 u}{\partial x^2}(x_i, t) \approx \frac{u(x_{i+1}, t) - 2u(x_i, t) + u(x_{i-1}, t)}{(\Delta x)^2}, \quad \Delta x = \frac{1}{n+1}.$$

763 This leads to the system of ODEs.

$$764 \quad \frac{d\tilde{\mathbf{u}}(t)}{dt} = \mathbf{A}\tilde{\mathbf{u}}(t),$$

766 where $\mathbf{A} \in \mathbb{R}^{n \times n}$ is the matrix

$$767 \quad \mathbf{A} := \frac{1}{(n+1)^2} \begin{pmatrix} -2 & 1 & 0 & \cdots & 0 \\ 1 & -2 & 1 & \cdots & 0 \\ 0 & 1 & -2 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & -2 \end{pmatrix},$$

773 and $\tilde{\mathbf{u}}(t) \in \mathbb{R}^n$ is the discretization of $u(\cdot, t)$ on the grid \mathcal{M} . Consequently, the solution to this
 774 system of ODEs can be expressed in terms of the matrix exponential $\tilde{\mathbf{u}}(t) = \exp(t\mathbf{A})\tilde{\mathbf{u}}(0)$.
 775

776 A.4 TV DISTANCE BOUND

778 **Theorem A.2.** *Let $\mathcal{P}_t, \mathcal{P}_{t+\delta}$ be the laws of $u_t, u_{t+\delta}$. Under the assumptions and with the notation
 779 from theorem 4.2,*

$$780 \quad d_{\text{TV}}(\mathcal{P}_t, \mathcal{P}_{t+\delta}) = \frac{1}{2} \int_{\mathbb{R}^n} |p_t(\mathbf{x}) - p_{t+\delta}(\mathbf{x})| d\mathbf{x}, \quad (2)$$

782 where $p_t, p_{t+\delta}$ are the densities of $\mathcal{P}_t, \mathcal{P}_{t+\delta}$, and

$$783 \quad d_{\text{TV}}(\mathcal{P}_t, \mathcal{P}_{t+\delta}) \leq \sqrt{\frac{1}{4} \left[\text{tr}(\Sigma_{t+\delta}^{-1} \Sigma_t) - n + (\Delta\mu)^T \Sigma_{t+\delta}^{-1} \Delta\mu + \log \frac{\det(\Sigma_{t+\delta})}{\det(\Sigma_t)} \right]}, \quad (3)$$

787 where

$$788 \quad \Delta\mu = (\mu_{t+\delta} - \mu_t) \quad \mu_t = \exp(t\mathbf{A})\mu_0 + \int_0^t \exp((t-s)\mathbf{A})\mathbf{r}(s)ds, \quad \Sigma_t = \exp(t\mathbf{A})\Sigma_0 \exp(t\mathbf{A}^T).$$

791 *Proof.* We know that $\mathcal{P}_t, \mathcal{P}_{t+\delta}$ admit densities from theorem 4.2. 2 follows by the definition of the
 793 TV distance between two distributions with densities $p_t, p_{t+\delta}$. For 3, we use Pinsker's inequality,
 794 which yields an upper bound on the TV distance by the Kullback–Leibler divergence

$$795 \quad d_{\text{TV}}(\mathcal{N}(\mu_1, \Sigma_1), \mathcal{N}(\mu_2, \Sigma_2)) \leq \sqrt{\frac{1}{2} D_{\text{KL}}(\mathcal{N}(\mu_1, \Sigma_1) \parallel \mathcal{N}(\mu_2, \Sigma_2))}.$$

798 The Kullback–Leibler divergence of two Gaussians above is well known, and given by

$$800 \quad \frac{1}{2} \left[\text{tr}(\Sigma_2^{-1} \Sigma_1) - n + (\mu_2 - \mu_1)^T \Sigma_2^{-1} (\mu_2 - \mu_1) + \log \frac{\det(\Sigma_2)}{\det(\Sigma_1)} \right].$$

802 Thus the claim follows by plugging in the calculated means and covariances. \square

803

804

805

806 A.5 DATA GENERATION

808 We generated multiple synthetic datasets from a general second order formulation for periodic PDEs,
 809 each discretized with finite-difference schemes. Initial conditions were sampled from Gaussian
 810 processes with covariance $(-\partial_x^2 + 25I)^{-2}$ to provide smooth but nontrivial trajectories. Each dataset

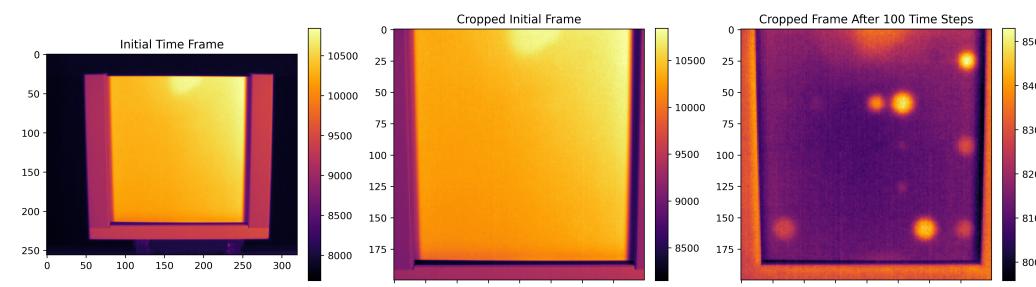
810 is stored as a compressed `.npz` file containing $(n_{\text{samples}}, n_t + 1, n_x)$ trajectories together with grid
811 and metadata. The PDE looks as follows

$$\begin{aligned} 812 \quad u_t(x, t) + au(x, t)_{xx} + bu(x, t)_x + cu(x, t) &= 0, & x \in (0, 1), t \geq 0, \\ 813 \quad u(0, t) = u(1, t) &= 0, & t \geq 0, \\ 814 \quad u(x, 0) = u_0(x), & x \in (0, 1). \end{aligned}$$

815 where $a, b, c \in \mathbb{R}$ are tuneable parameters. We sampled data from all combinations of $a \in \{ -0.005, -0.0075, -0.01 \}$, $b \in \{ -0.5, 0, 0.5 \}$, and $c \in \{ -0.5, 0, 0.5 \}$. Note that we focused
816 on negative values for a , as positive values tend to smoothen the trajectory over time which leads
817 to trivial coverage even for naïve CP. We sampled 21 time steps (incl. initial condition) between 0
818 and 1.
819

820 A.6 REAL-WORLD EXAMPLE

821 The dataset of Wei et al. (2023) contains 19 samples of 2D surfaces measured over 1810 time steps.
822 We allocate 8 samples for training a Fourier Neural Operator Li et al. (2021), 1 for validation, 7
823 for calibration, and 3 for testing. Calibration begins at the end of the heating process, and testing
824 proceeds over the subsequent time steps. We crop the borders of the measurement data, since they
825 correspond to non-heated background regions (see figure 4). We use outlier clipping and a Gaussian
826 filter to get rid of the measurement noise, and normalize the data with values of the first time step.
827 Because the dataset is small, we divide each surface into square patches of size 5×5 and treat them
828 as identically distributed samples (to mitigate dependency between the patches, we leave an empty
829 space of size 5 to all neighboring patches). While this induces some dependence between patches,
830 and the variations in the metal plates (e.g., differently drilled backside holes) introduce slight
831 distribution shifts, both effects are minor for a real-world measurement dataset and have negligible
832 impact on CP performance. As the system cools, both the data and prediction residuals smooth out
833 over time, causing many CP methods to overcover rather than undercover. We show the results in
834 figure 5. Weighted CP, which only detects that the distribution has changed, still produces infinite
835 intervals. Despite this limitation and the relative advantage of standard CP in this specific setting,
836 our experiment shows that the weighted CP approach remains practically applicable.
837



838 Figure 4: Exemplary frame from the real-world heat pulsed-thermography dataset. The left image
839 shows a raw data frame from the beginning of the cooling period. The middle image shows how we
840 crop the frames to remove the background. The right image shows how the data looks after 100
841 time steps.

842 A.7 ADDITIONAL RESULTS ON SYNTHETIC DATA

843 We show results for remaining b -values in figures 6 and 7. For $b = -0.5$, we report the remaining
844 numbers in tables 2 and 3. In Figure 8, we see an example of CP bands where naïve CP and LSCI
845 undercover and WCP remains full coverage.
846

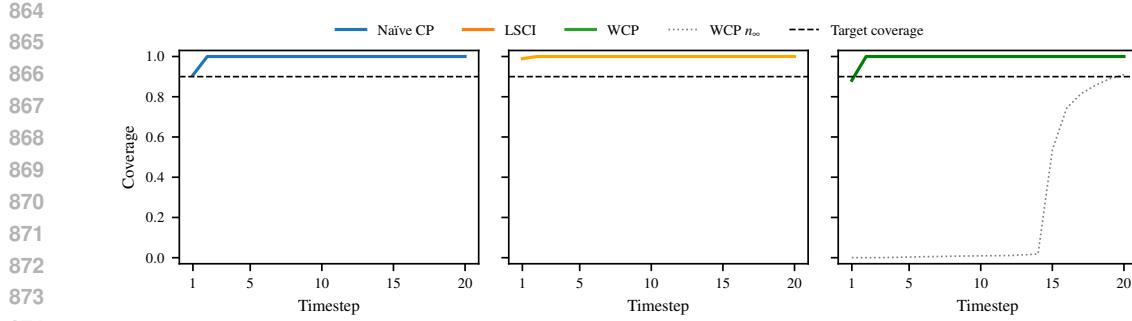


Figure 5: Coverages of the CP methods on the heat pulsed-thermography dataset. As the data smoothens out over time, coverages become trivial at some point.

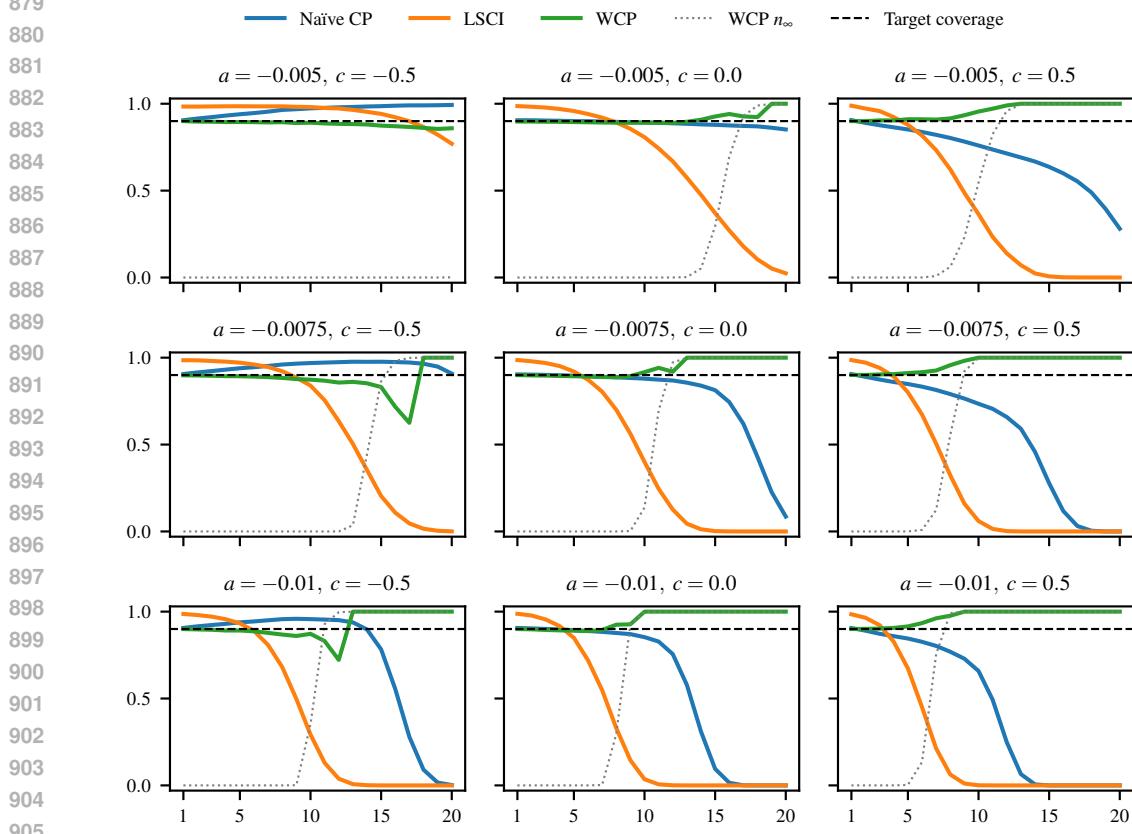


Figure 6: Mean coverages for varying a, c and fixed $b = 0$ across increasing prediction horizon. We omit coverages when infinite conformal bands were reported (coverage of 1 would hold trivially). The 90% coverage target and the fraction of samples with infinite bandwidth n_∞ are marked as black dashed line and gray dotted line, respectively.

A.8 OTHER INITIAL DISTRIBUTIONS

To show that our method also works when sampling data from other distributions of the location-scale family, we include coverage statistics when using the Laplace and the logistic distribution for the initial values (only for $b = -0.5$, to keep this outlook concise). The figure for Laplace can be found in 9 and for logistic in 10.

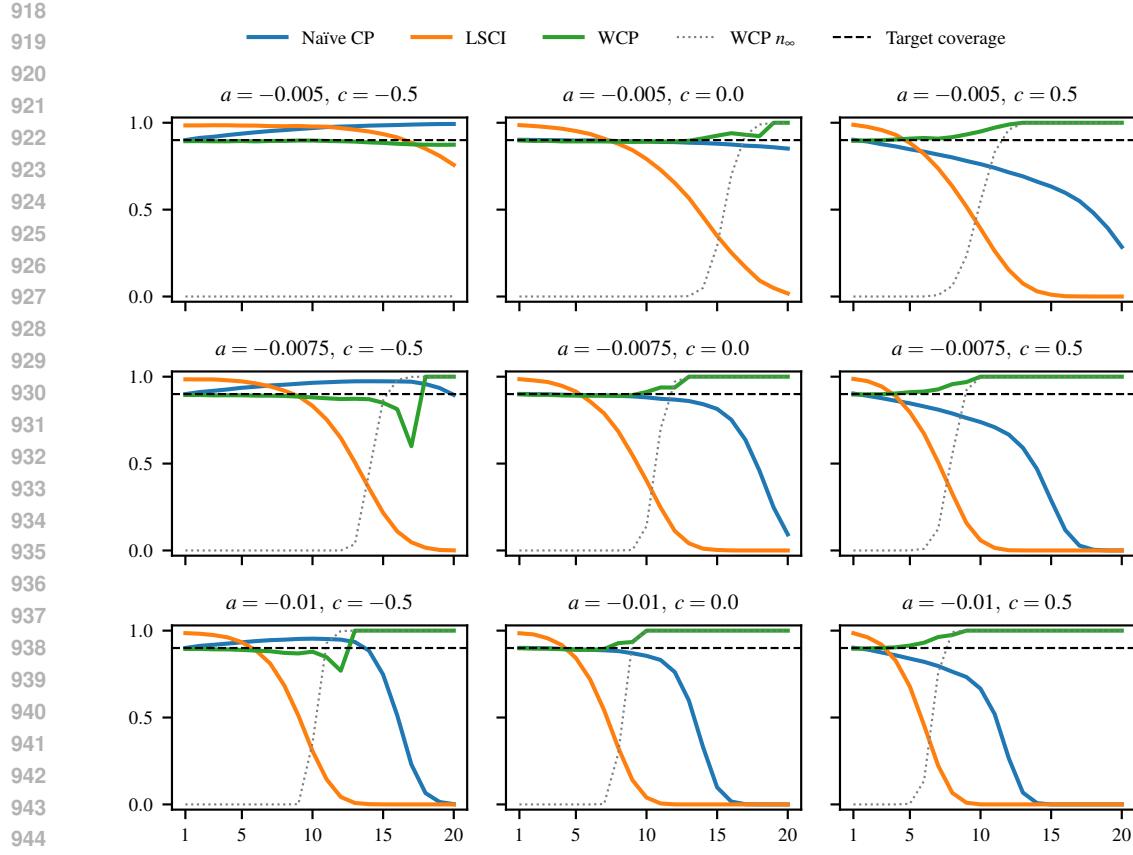


Figure 7: Mean coverages for varying a, c and fixed $b = 0.5$ across increasing prediction horizon. We omit coverages when infinite conformal bands were reported (coverage of 1 would hold trivially). The 90% coverage target and the fraction of samples with infinite bandwidth n_∞ are marked as black dashed line and gray dotted line, respectively.

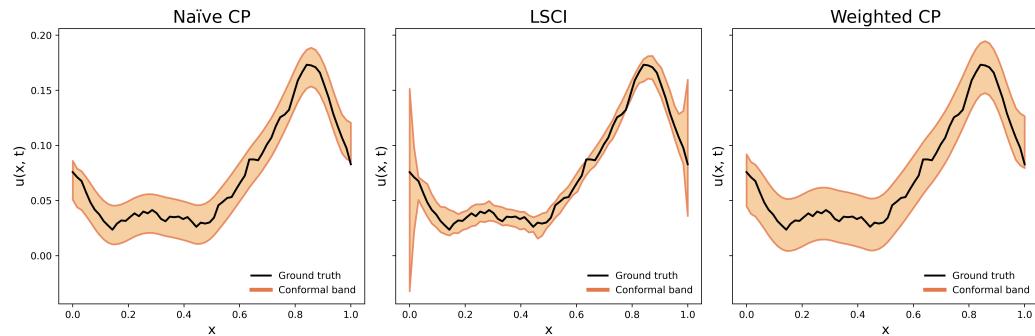


Figure 8: Exemplary coverage plots for the three methods naïve CP (left), LSCI (middle), and WCP (right) for a sample at time step 7 with parameters $a = -0.005, b = 0.5, c = 0.5$. In this case, only WCP fully covers the the trajectory.

A.9 USE OF LARGE LANGUAGE MODELS

Large language models (LLMs) helped in the creation and execution of this project. They assisted with improving the clarity and readability of the manuscript, suggesting alternative phrasings, providing feedback on mathematical arguments, and offering ideas during the research and coding process. All research contributions, results, and final formulations were verified manually.

972
 973
 974
 975
 976
 977
 978
 979
 980
 981
 982
 983
 984

985 Table 2: Mean coverages and bandwidths over 5000 sampled initial conditions for varying a and
 986 fixed $b = -0.5$, $c = 0$. For WCP, we also report the fraction of samples where infinite bands were
 987 reported (n_∞) to maintain coverage guarantees. The gray font is chosen for better readability.

988
 989
 990
 991
 992
 993
 994
 995
 996
 997
 998

			Timestep				
			1	5	10	15	20
$a = -0.005$	Naïve CP	Coverage	0.9	0.9	0.9	0.88	0.85
		Bandwidth	0.03	0.03	0.03	0.03	0.03
	LSCI	Coverage	0.99	0.96	0.82	0.39	0.03
		Bandwidth	0.02	0.02	0.02	0.02	0.02
	WCP (Ours)	Coverage	0.9	0.9	0.9	0.93	1.0
		Bandwidth	0.03	0.03	0.03	0.04	0.04
$a = -0.0075$	Naïve CP	Coverage	0.9	0.9	0.88	0.8	0.08
		Bandwidth	0.03	0.03	0.03	0.03	0.03
	LSCI	Coverage	0.99	0.92	0.43	0.0	0.0
		Bandwidth	0.02	0.02	0.02	0.02	0.02
	WCP (Ours)	Coverage	0.9	0.9	0.92	1.0	1.0
		Bandwidth	0.03	0.03	0.04	∞	∞
$a = -0.01$	Naïve CP	Coverage	0.91	0.9	0.86	0.09	0.0
		Bandwidth	0.03	0.03	0.03	0.03	0.03
	LSCI	Coverage	0.99	0.86	0.04	0.0	0.0
		Bandwidth	0.02	0.02	0.02	0.02	0.02
	WCP (Ours)	Coverage	0.9	0.9	1.0	1.0	1.0
		Bandwidth	0.03	0.03	0.07	∞	∞
		n_∞	0.0%	0.0%	99.9%	100%	100%

1013
 1014
 1015
 1016
 1017
 1018
 1019
 1020
 1021
 1022
 1023
 1024
 1025

1026
 1027
 1028
 1029
 1030
 1031
 1032
 1033
 1034
 1035
 1036
 1037
 1038

1039 Table 3: Mean coverages and bandwidths over 5000 sampled initial conditions for varying a and
 1040 fixed $b = -0.5$, $c = 0.5$. For WCP, we also report the fraction of samples where infinite bands were
 1041 reported (n_∞) to maintain coverage guarantees. The gray font is chosen for better readability.

1042
 1043
 1044
 1045
 1046
 1047
 1048
 1049
 1050
 1051
 1052
 1053
 1054
 1055
 1056
 1057
 1058
 1059
 1060
 1061
 1062
 1063
 1064
 1065
 1066
 1067
 1068
 1069
 1070
 1071
 1072
 1073
 1074
 1075
 1076
 1077
 1078
 1079

		Timestep					
			1	5	10	15	20
$a = -0.005$	Naïve CP	Coverage	0.9	0.85	0.76	0.63	0.27
		Bandwidth	0.03	0.03	0.03	0.03	0.03
	LSCI	Coverage	0.99	0.89	0.4	0.01	0.0
		Bandwidth	0.02	0.02	0.02	0.02	0.02
	WCP (Ours)	Coverage	0.9	0.91	0.95	1.0	1.0
		Bandwidth	0.03	0.04	0.05	0.07	∞
$a = -0.0075$	Naïve CP	Coverage	0.91	0.85	0.74	0.26	0.0
		Bandwidth	0.03	0.03	0.03	0.03	0.03
	LSCI	Coverage	0.99	0.82	0.07	0.0	0.0
		Bandwidth	0.02	0.02	0.02	0.02	0.02
	WCP (Ours)	Coverage	0.9	0.91	1.0	1.0	1.0
		Bandwidth	0.03	0.04	0.06	∞	∞
$a = -0.01$	Naïve CP	Coverage	0.9	0.85	0.65	0.0	0.0
		Bandwidth	0.03	0.03	0.03	0.03	0.03
	LSCI	Coverage	0.99	0.71	0.0	0.0	0.0
		Bandwidth	0.02	0.02	0.02	0.02	0.02
	WCP (Ours)	Coverage	0.9	0.91	1.0	1.0	1.0
		Bandwidth	0.03	0.04	∞	∞	∞
		n_∞	0.0%	0.0%	99.7%	100%	100%

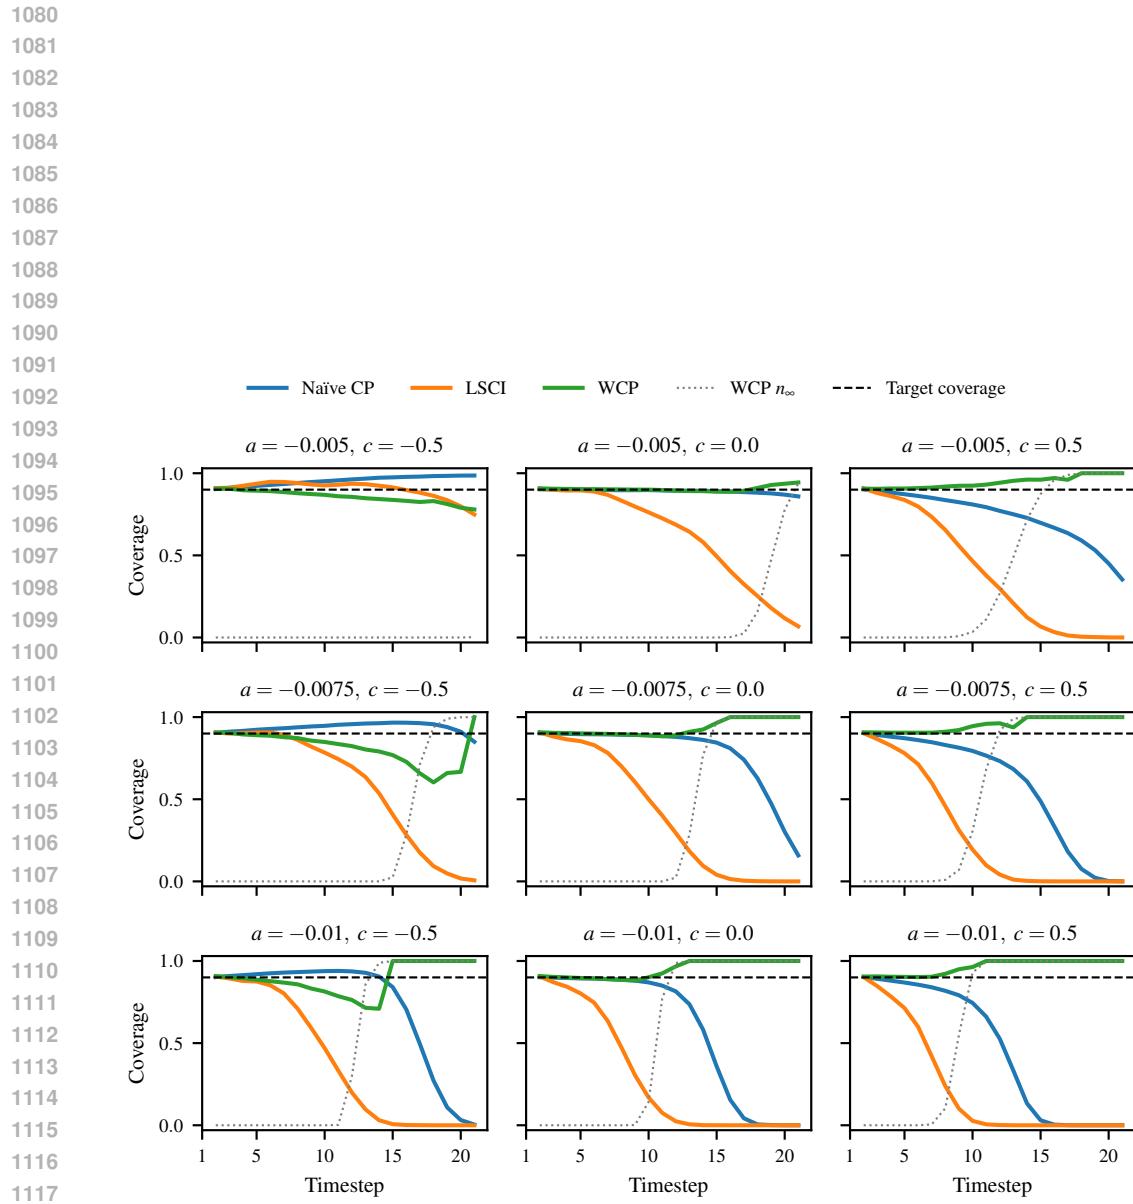
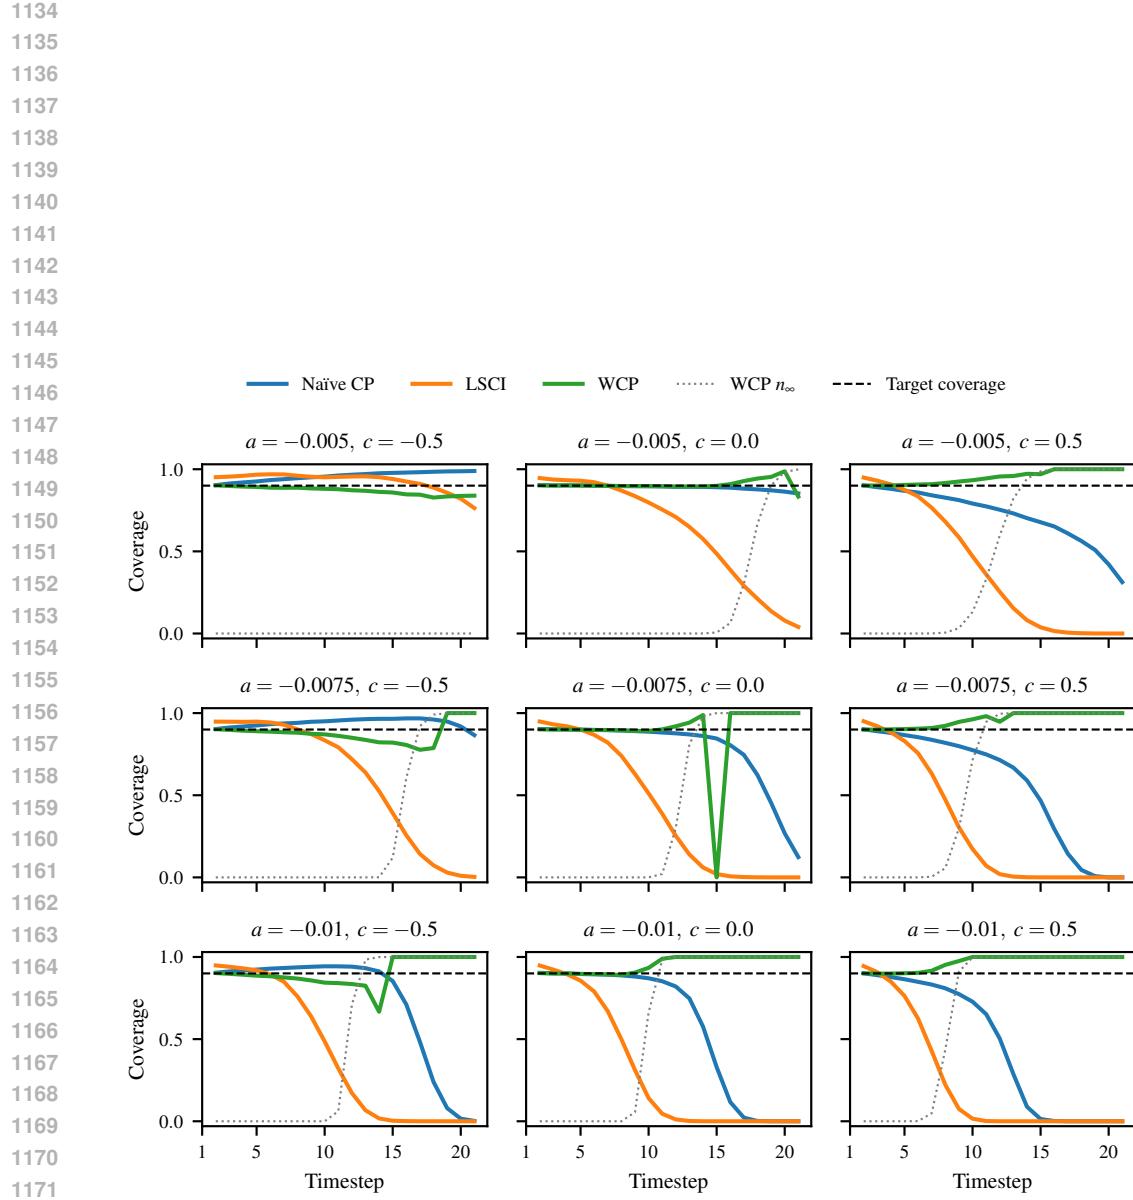


Figure 9: Mean coverages when sampling from a Laplacian initial distribution for varying a, c and fixed $b = -0.5$ across increasing prediction horizon. We omit coverages when infinite conformal bands were reported (coverage of 1 would hold trivially). The 90% coverage target and the fraction of samples with infinite bandwidth n_∞ are marked as black dashed line and gray dotted line, respectively.

1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133



1173
 1174
 1175
 1176
 1177
 1178
 1179
 1180
 1181
 1182
 1183
 1184
 1185
 1186
 1187