CERTIFYING SIMULATION-BASED INFERENCE FOR NONLINEAR SCIENTIFIC MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Simulation-based inference (SBI) learns information for posterior inference from simulator outputs when likelihood evaluation is unavailable or expensive. In scientific applications, however, statistical accuracy of the learned posterior does not guarantee that every retained parameter value defines a mathematically well-posed continuous model, nor that an approximate solver is close to that model's solution. We introduce *certified SBI*, an estimator-agnostic layer that conditions a learned posterior on parameters passing a deterministic model-validity certificate. The framework covers both time-dependent models described by fixed-point maps and steady-state models described by residual equations. For certified parameters, we prove sample-wise existence and uniqueness, residual-to-error bounds for approximate solvers, stability of posterior conditioning, and end-to-end bounds that propagate certification and posterior-approximation errors to downstream posterior summaries and Bayes decisions. The certification interface is compatible with neural posterior, likelihood, or ratio estimation and with classical or learned PDE solvers. A controlled nonlinear inverse problem demonstrates the resulting validity–coverage–cost trade-off.

1 INTRODUCTION

Simulation-based inference (SBI) enables Bayesian inference when a scientific simulator can generate observations but its likelihood is unavailable or too expensive to evaluate (Cranmer et al., 2020; Lueckmann et al., 2021). The modern SBI toolbox typically includes neural posterior estimation (NPE), which learns a conditional distribution over parameters; neural likelihood estimation (NLE), which learns the simulator likelihood; and neural ratio estimation (NRE), which learns likelihood-to-evidence or related density ratios (Papamakarios et al., 2019; Greenberg et al., 2019; Hermans et al., 2020).

Most SBI guarantees and diagnostics focus on the statistical layer: whether the learned posterior approximates the intended Bayesian target, whether inference is calibrated under repeated simulation, and how efficiently simulations are allocated (Talts et al., 2018; Krouglova et al., 2026). This leaves a distinct reliability gap when the simulator itself is a nonlinear scientific model. Statistical calibration does not establish that an individual posterior draw lies in a regime where the underlying continuous equation has a well-defined and stable solution. Bayesian inverse-problem theory typically assumes that the parameter-to-observation map is well defined (Stuart, 2010; Dashti and Stuart, 2017); standard SBI then treats that map as a simulator. Thus posterior reliability and continuous-model validity remain logically separate questions.

We address this gap with *certified SBI*, a modular layer that makes continuous-model validity explicit before posterior samples are used for prediction, design, or control. The idea is simple: train an SBI estimator in the usual way; propose a posterior draw; evaluate a deterministic certificate derived from the continuous scientific model; and retain the draw only when the certificate succeeds. Rejection does not mean that the parameter is physically impossible or that its solution is singular. It means only that the selected mathematical theorem does not validate that draw. Consequently, the method returns both a certified posterior and the posterior mass not covered by the chosen theorem.

After an SBI estimator has learned an approximate posterior $Q_{N,y}$ from N simulations and observed data y, a model-specific predicate $C(\theta)$ checks whether a proposed parameter θ satisfies the assumptions of a deterministic existence-and-stability theorem. Passing draws form a conditional posterior

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

$Q^{\Gamma}_{N,y}$ on the certified set Γ. The output is therefore not only a distribution over plausible parameters: each retained draw comes with a specified continuous solution class and, when an approximate solver is used, a quantitative state-error guarantee. The same accept/reject layer can be placed after NPE, NLE, or NRE without changing their training objectives. This separates posterior approximation from continuous-model verification. A more accurate SBI estimator reduces statistical approximation error; a better numerical solver reduces state-approximation error; and a sharper analytical theorem enlarges the set of draws that can be certified. These improvements are composable rather than mutually dependent. At the same time, certification is not free: conditioning can alter posterior summaries and decisions, and certificate evaluation adds computation. Our theory therefore quantifies validity, posterior distortion, decision error, and computational cost within the same framework.

Contributions. Our contributions are:

- **An estimator-agnostic sample-level validity layer.** We convert deterministic existence, uniqueness, and stability conditions into a measurable posterior operator. In contrast to methods that improve the learned density or simulation allocation, certified SBI addresses a different question: whether a retained draw corresponds to a mathematically controlled continuous model instance.
- **A common verification interface for classical and learned solvers.** We develop contraction- and residual-stability certificates that turn continuous defects into state-error bounds. The verification theorem is independent of whether the approximate state is produced by a mesh-based solver, a PINN, or a neural operator, provided a rigorous bound on the corresponding continuous defect is available.
- **End-to-end statistical consequences of certification.** We prove stability of the conditioned posterior, bounds for bounded posterior functionals, and excess Bayes decision risk. Here a reported prediction is any bounded posterior summary used downstream, while a Bayes decision is an action chosen to minimize posterior expected loss. The bounds therefore control how much a scientific summary or decision can change because uncertified mass was removed.
- **An explicit validity–coverage–cost trade-off.** We connect retained posterior mass to rejection complexity and certificate cost, making the computational price of certification explicit. The experiments quantify retained posterior mass, posterior shift, certificate-threshold sensitivity, and rejection overhead rather than reporting validity as a binary label.
- **Reusable instantiations across nonlinear scientific models.** The same posterior layer is instantiated with distinct deterministic theories for reaction–diffusion, incompressible flow, nonlinear conductivity, and steady-state nonlinear diffusion, illustrating that the statistical mechanism does not depend on one PDE family.

Organization. Section 2 defines the certificate interfaces and certified posterior. Section 3 establishes state-, posterior-, decision-, and cost-level guarantees. Section 4 gives the certified SBI sampling algorithm and its diagnostics. Section 5 presents scientific-model realizations and experiments. Section 6 reviews related work, Section 7 discusses limitations, and Section 8 concludes. Complete proofs, numerical-verification details, and additional model instantiations are provided in the appendix.

2 Certification Framework

We formalize certificates at the level of the continuous scientific model, independently of the posterior estimator.

2.1 Scientific Model

Let $\Theta \subseteq \mathbb{R}^{d_\theta}$ be the parameter space and $\mathcal{B}(\Theta)$ its Borel σ-algebra. Let $(\mathcal{X}, \|\cdot\|_{\mathcal{X}})$ be a Banach state space, and let $\mathcal{X}^*$ denote its continuous dual. For time-dependent models, $\mathcal{U}$ denotes a Banach data space and the initial condition is written $u_0 \in \mathcal{U}$; boundary or forcing data may be included in the same data variable. These spaces are determined by the deterministic existence and stability theorem used for certification, rather than by the SBI architecture.

For a fixed $\theta \in \Theta$, the *fixed-point representation*

$$u = T_\theta(u), \qquad T_\theta : \mathcal{X} \to \mathcal{X}, \tag{1}$$

is natural when the model admits a state-update map, as in many diffusion, flow, and wave equations. The parameter θ may encode physical coefficients, forcing, initial data, or other inferred quantities. For models specified more naturally through a governing balance, we use

$$\mathcal{R}_\theta(u) = 0, \qquad \mathcal{R}_\theta : \mathcal{X} \to \mathcal{Y}, \tag{2}$$

where $\mathcal{Y}$ is a Banach *residual space* whose norm measures violation of the continuous equation. In weak formulations we typically take $\mathcal{Y} = \mathcal{X}^*$. This representation covers, for example, steady diffusion and conductivity equations and other variational models, while learned or classical numerical solvers can be assessed by evaluating the same continuous residual at their approximate states.

For $R > 0$, write $\overline{B}_{\mathcal{X}}(0, R) := \{u \in \mathcal{X} : \|u\|_{\mathcal{X}} \le R\}$ for the closed ball of candidate states. By *well-posedness of the scientific state problem* we mean existence, uniqueness in a specified solution class, and continuous dependence on the model data in the corresponding topology. The precise solution-class and stability definitions used here are given in Appendix A.1.

2.2 Certification Criteria

Fix $\theta \in \Theta$, hence one concrete scientific model. The certificate verifies quantitative conditions for the corresponding equation (1) or (2).

Route I: contraction certificate. Choose a radius $R_\theta > 0$ and a contraction factor $\lambda_\theta \in [0, 1)$ such that

$$\|T_\theta(0)\|_{\mathcal{X}} + \sup_{u \in \overline{B}_{\mathcal{X}}(0, R_\theta)} \|T_\theta(u) - T_\theta(0)\|_{\mathcal{X}} \le R_\theta, \tag{3}$$

$$\|T_\theta(u) - T_\theta(v)\|_{\mathcal{X}} \le \lambda_\theta \|u - v\|_{\mathcal{X}}, \qquad u, v \in \overline{B}_{\mathcal{X}}(0, R_\theta). \tag{4}$$

The first inequality verifies invariance of the candidate-state region $\overline{B}_{\mathcal{X}}(0, R_\theta)$, while the second makes T_θ a strict contraction there. By Banach's fixed-point theorem, these conditions yield existence and uniqueness of a fixed point in $\overline{B}_{\mathcal{X}}(0, R_\theta)$; see Theorem 1 and Appendix D.

Route II: residual-stability certificate. For a residual model (2), we use the dual pairing $\langle \cdot, \cdot \rangle_{\mathcal{X}^*, \mathcal{X}} : \mathcal{X}^* \times \mathcal{X} \to \mathbb{R}$. The certificate supplies a quantitative stability modulus $\omega_\theta : [0, \infty) \to [0, \infty)$ with $\omega_\theta(0) = 0$ and $\omega_\theta(s) > 0$ for $s > 0$, satisfying, for all candidate states $u, v \in \mathcal{X}$ in the specified solution class,

$$\langle \mathcal{R}_\theta(u) - \mathcal{R}_\theta(v), u - v \rangle_{\mathcal{X}^*, \mathcal{X}} \ge \omega_\theta(\|u - v\|_{\mathcal{X}}). \tag{5}$$

A positive modulus makes state differences detectable through the continuous residual. Standard monotone-operator assumptions provide existence and uniqueness, while quantitative monotonicity yields stability moduli of this form (Browder, 1963; Minty, 1962; Showalter, 1997); see Theorem 2 and Appendix D.

Define the *certificate indicator* $C : \Theta \to \{0, 1\}$ by setting $C(\theta) = 1$ exactly when the assumptions of either the Route I certificate in Theorem 1 or the Route II certificate in Theorem 2 are verified for θ, and set $C(\theta) = 0$ otherwise. We assume that C is $\mathcal{B}(\Theta)$-measurable and define $\Gamma := \{\theta \in \Theta : C(\theta) = 1\}$. This measurability requirement is mild in the finite-dimensional parameter settings considered here. Concrete sufficient conditions are given in Appendix D.5. In implementations, C is often obtained by thresholding one or more real-valued certificate quantities, such as a contraction factor λ_θ or a verified stability bound. These quantities encode the sufficient inequalities required by the selected deterministic theorem.

Definition 1 (Certification operator)**.** *Let $\mathcal{P}(\Theta)$ denote the probability measures on $(\Theta, \mathcal{B}(\Theta))$. For any $Q \in \mathcal{P}(\Theta)$ with $Q(\Gamma) > 0$, define*

$$\mathrm{Cert}_\Gamma(Q)(A) := \frac{Q(A \cap \Gamma)}{Q(\Gamma)}, \qquad A \in \mathcal{B}(\Theta). \tag{6}$$

Equation (6) is ordinary conditioning on the measurable certified set Γ.

Definition 2 (Learned and target certified posteriors)**.** *Let $Q_{N,y} \in \mathcal{P}(\Theta)$ be the posterior approximation learned from N simulations after observing data y, and let $\Pi_y \in \mathcal{P}(\Theta)$ be the target Bayesian posterior under the stated simulator and observation model. If $Q_{N,y}(\Gamma) > 0$ and $\Pi_y(\Gamma) > 0$, define*

$$Q^\Gamma_{N,y} := \mathrm{Cert}_\Gamma(Q_{N,y}), \qquad \Pi^\Gamma_y := \mathrm{Cert}_\Gamma(\Pi_y).$$

Thus $Q^{\Gamma}_{N,y}$ is the distribution sampled by certified SBI, whereas Π^{Γ}_{y} is its ideal conditional target. Well-definedness, and normalization of the conditioning operator are verified in Appendix D.5.

3 Guarantees for Certified Inference

We now establish the three guarantees needed by the method: accepted parameters have well-defined scientific states, approximate solver outputs can be related to those states, and conditioning has controlled statistical and computational effects.

3.1 Sample-Wise Well-Posedness

Theorem 1 (Contraction-certified well-posedness)**.** *Let $\theta \in \Gamma$ satisfy* (3)–(4)*. Then the fixed-point equation $u = T_\theta(u)$ in* (1) *has a unique solution $u^\theta \in \overline{B}_{\mathcal{X}}(0, R_\theta)$. Consequently, every sample accepted through Route I corresponds to one uniquely defined continuous state in the certified state region $\overline{B}_{\mathcal{X}}(0, R_\theta)$.*

The proof is in Appendix F. We next give the corresponding existence-and-uniqueness result for the residual route. A concrete Route II instance writes $\mathcal{R}_\theta(u) = \mathcal{A}_\theta(u) - f_\theta$, where $\mathcal{A}_\theta : \mathcal{X} \to \mathcal{X}^*$ is the model operator and $f_\theta \in \mathcal{X}^*$ represents forcing, boundary data after lifting, or another source term.

Theorem 2 (Residual-certified well-posedness)**.** *Let $\mathcal{X}$ be reflexive and let $\theta \in \Gamma$. Suppose*

$$\mathcal{R}_\theta(u) = \mathcal{A}_\theta(u) - f_\theta, \qquad \mathcal{A}_\theta : \mathcal{X} \to \mathcal{X}^*, \quad f_\theta \in \mathcal{X}^*,$$

where $\mathcal{A}_\theta$ is bounded on bounded sets, hemicontinuous, and coercive. Assume the Route II stability condition (5) *holds for all $u, v \in \mathcal{X}$, with $\omega_\theta(s) > 0$ for $s > 0$. Then $\mathcal{R}_\theta(u) = 0$, equivalently $\mathcal{A}_\theta(u) = f_\theta$, has a unique solution $u^\theta \in \mathcal{X}$.*

Because f_θ cancels in differences, $\mathcal{R}_\theta(u) - \mathcal{R}_\theta(v) = \mathcal{A}_\theta(u) - \mathcal{A}_\theta(v)$; hence the stability condition (5) is equivalently a quantitative monotonicity condition on $\mathcal{A}_\theta$. Existence follows from the Browder–Minty theorem and uniqueness from the positive stability modulus (Browder, 1963; Minty, 1962; Showalter, 1997). Concrete choices of $\mathcal{A}_\theta$ and f_θ for the p-Laplacian, nonlinear conductivity, and reaction–diffusion equilibria are collected in Appendix D.

3.2 A Posteriori Solver Certification

For Route I, define the fixed-point defect operator

$$\mathcal{E}_\theta : \overline{B}_{\mathcal{X}}(0, R_\theta) \to \mathcal{X}, \qquad \mathcal{E}_\theta(v) := v - T_\theta(v). \tag{7}$$

Thus $\mathcal{E}_\theta(v) = 0$ exactly at a fixed point. Throughout this paper, an *approximate state* means any numerical or learned solver output v^θ belonging to the stated function space; it is not assumed to satisfy the continuous equation exactly.

Theorem 3 (Residual-to-error bound: contraction route)**.** *Let $\theta \in \Gamma$ satisfy* (3)–(4)*, and let $u^\theta \in \overline{B}_{\mathcal{X}}(0, R_\theta)$ be the unique solution from Theorem 1. For any approximate state $v^\theta \in \overline{B}_{\mathcal{X}}(0, R_\theta)$,*

$$\left\|u^\theta - v^\theta\right\|_{\mathcal{X}} \le (1 - \lambda_\theta)^{-1} \left\|\mathcal{E}_\theta(v^\theta)\right\|_{\mathcal{X}}. \tag{8}$$

Theorem 3 shows how a continuous solver defect translates into a state error. The gap $1 - \lambda_\theta > 0$ from the contraction threshold $\lambda_\theta = 1$ controls this error amplification, with the bound deteriorating as $\lambda_\theta \uparrow 1$. Because the guarantee is attached to the continuous equation rather than to one solver architecture, a mesh solver can be replaced by a PINN or neural operator whenever its output admits a rigorous bound on the same continuous defect. Solver and posterior-estimator improvements can therefore be developed independently of the mathematical certificate.

Theorem 4 (Residual-to-error bound: stability route)**.** *Let $\theta \in \Gamma$ satisfy the Route II assumptions of Theorem 2, let $u^\theta \in \mathcal{X}$ be its unique solution of $\mathcal{R}_\theta(u^\theta) = 0$, and let $v^\theta \in \mathcal{X}$ be an approximate state. Assume* (5)*. Then, with $r^\theta := \mathcal{R}_\theta(v^\theta) \in \mathcal{X}^*$,*

$$\omega_\theta(\left\|u^\theta - v^\theta\right\|_{\mathcal{X}}) \le \left\|r^\theta\right\|_{\mathcal{X}^*} \left\|u^\theta - v^\theta\right\|_{\mathcal{X}}. \tag{9}$$

If, in addition, $\omega_\theta(s) \ge m_\theta s^p$ for all $s \ge 0$ for some $p > 1$ and $m_\theta > 0$, then

$$\left\|u^\theta - v^\theta\right\|_{\mathcal{X}} \le m_\theta^{-1/(p-1)} \left\|r^\theta\right\|_{\mathcal{X}^*}^{1/(p-1)}. \tag{10}$$

Strong monotonicity corresponds to the common case $p = 2$. The proof is in Appendix F. Theorem 4 gives the analogous guarantee without a fixed-point representation: a small continuous residual becomes a state-error bound through the stability modulus. This makes Route II useful for steady-state conductivity, p-Laplacian, and other quasilinear inverse problems, where a Duhamel-style state update is unnatural.

3.3 Posterior Stability, Decisions, and Cost

The state-level results above answer whether an accepted parameter and its computed state are mathematically controlled. We next quantify how conditioning changes the learned posterior, its expectations, downstream decisions, and sampling cost. We use total variation

$$\mathrm{TV} : \mathcal{P}(\Theta)^2 \to [0,1], \qquad \mathrm{TV}(P,Q) := \sup_{A \in \mathcal{B}(\Theta)} |P(A) - Q(A)|.$$

Let $D : \mathcal{P}(\Theta)^2 \to [0,\infty]$ be any divergence for which a nondecreasing calibration function $\psi_D : [0,\infty] \to [0,1]$ satisfies

$$\mathrm{TV}(P,Q) \le \psi_D(D(Q\|P)). \tag{11}$$

This includes Kullback–Leibler divergence and many f-divergences for which Pinsker-type inequalities are available (Reid and Williamson, 2011; Gilardoni, 2010; Sason and Verdú, 2016). For KL divergence, Pinsker's inequality allows the choice $\psi_D(t) = \sqrt{t/2}$. The precise f-divergence formulation, including singular measures, is given in Appendix E.

Theorem 5 (Certified posterior stability)**.** *Fix a measurable observation space $\mathcal{Y}_{\mathrm{obs}}$, observed data $y \in \mathcal{Y}_{\mathrm{obs}}$, a measurable certified set Γ, the target posterior Π_y, and an SBI approximation $Q_{N,y}$ trained from N simulations. Assume $\Pi_y(\Gamma) \ge \alpha > 0$ and $D(Q_{N,y}\|\Pi_y) \le \varepsilon_N < \infty$. Set $r_N := \psi_D(\varepsilon_N)$. If $r_N < \alpha$, then*

$$\mathrm{TV}(Q^{\Gamma}_{N,y}, \Pi^{\Gamma}_{y}) \le \frac{2r_N}{\alpha - r_N}, \qquad Q_{N,y}(\Gamma) \ge \alpha - r_N. \tag{12}$$

Consequently, independent rejection sampling uses at most $(\alpha - r_N)^{-1}$ proposals in expectation per accepted draw.

Theorem 5 is conditional on an upstream SBI approximation bound $D(Q_{N,y}\|\Pi_y) \le \varepsilon_N$. When such a bound is available, it converts pre-certification posterior error into a post-certification guarantee, while $\alpha - r_N$ also controls rejection cost. Thus retained posterior mass is not merely a reporting statistic: it links the statistical stability of conditioning to its expected sampling cost.

To quantify the effect of certification on downstream posterior summaries and decisions, we use the following notation. For $P \in \mathcal{P}(\Theta)$ and P-integrable measurable $h : \Theta \to \mathbb{R}$, let $\mathbb{E}_P[h] := \int_\Theta h(\theta)\, P(d\theta)$ for bounded h, let $\|h\|_\infty := \sup_{\theta\in\Theta} |h(\theta)|$. Let $(\mathcal{A}, \Sigma_{\mathcal{A}})$ be a measurable decision space and $\ell : \mathcal{A} \times \Theta \to [0, L_{\max}]$ a measurable bounded loss, where $\ell(a,\theta)$ is the cost of action a at parameter θ. Its posterior risk is $R_P(a) := \mathbb{E}_P[\ell(a,\theta)]$.

Theorem 6 (End-to-end validity–accuracy bound)**.** *Under the assumptions of Theorem 5, let $\beta := \Pi_y(\Gamma)$ and $\Delta_N := (1-\beta) + 2r_N/(\beta - r_N)$. Then*

$$\mathrm{TV}(Q^{\Gamma}_{N,y}, \Pi_y) \le \Delta_N. \tag{13}$$

For every bounded measurable quantity h,

$$\left| \mathbb{E}_{Q^{\Gamma}_{N,y}} h - \mathbb{E}_{\Pi_y} h \right| \le 2\,\|h\|_\infty\, \Delta_N. \tag{14}$$

Moreover, suppose that the certified-posterior risk $a \mapsto R_{Q^{\Gamma}_{N,y}}(a)$ attains its minimum over $\mathcal{A}$, and let $\widehat{a} \in \arg\min_{a\in\mathcal{A}} R_{Q^{\Gamma}_{N,y}}(a)$ be a minimizing action. Then its excess risk relative to the optimal risk under the target posterior Π_y satisfies

$$R_{\Pi_y}(\widehat{a}) - \inf_{a\in\mathcal{A}} R_{\Pi_y}(a) \le 4L_{\max}\Delta_N.$$

Theorem 6 quantifies how certification and posterior approximation affect quantities ultimately computed from the posterior. The term $1 - \beta$ represents exact-posterior mass excluded by the

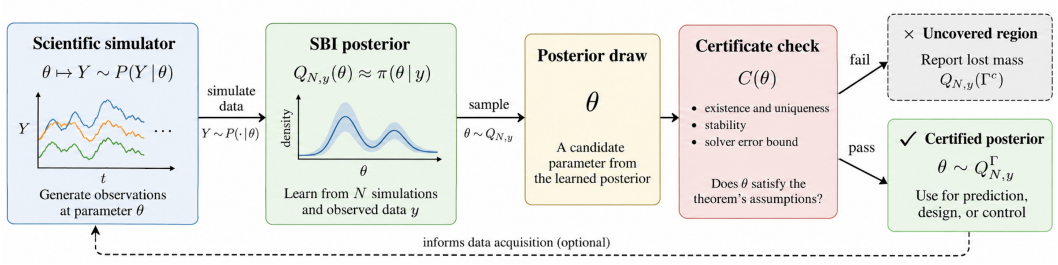


Figure 1: **Certified SBI workflow.** A scientific simulator generates observations used to learn an approximate posterior $Q_{N,y}$. For each candidate draw $\theta \sim Q_{N,y}$, the certificate $C(\theta)$ verifies whether the assumptions required for continuous-model well-posedness, stability, and solver-error control are satisfied. Passing draws define the certified posterior $Q^{\Gamma}_{N,y}$ for downstream prediction, design, or control, while failing draws quantify posterior mass not covered by the selected mathematical certificate. The dashed arrow is schematic and indicates a possible adaptive extension that is not used in the present experiments.

certificate, while $2r_N/(\beta - r_N)$ propagates the upstream posterior-approximation error through conditioning. The bound on $\left|\mathbb{E}_{Q^{\Gamma}_{N,y}}[h] - \mathbb{E}_{\Pi_y}[h]\right|$ controls bounded posterior summaries, while the decision bound controls the additional expected loss incurred when an action is selected using the certified posterior rather than the target posterior. This separation lets the experiments distinguish limited mathematical coverage from statistical approximation. Appendix E extends the result to integral probability metrics and more structured classes of test functions.

4 Certified Simulation-Based Inference

We now turn the preceding guarantees into an implementable inference procedure and connect each theorem to an observable diagnostic. Let $(\mathcal{Y}_{\text{obs}}, \Sigma_{\text{obs}})$ be a measurable observation space and let $\mathcal{G} : \Theta \to \mathcal{Y}_{\text{obs}}$ be a measurable parameter-to-observation map. Thus, for each parameter $\theta \in \Theta$, $\mathcal{G}(\theta)$ denotes the observation predicted by the corresponding continuous scientific model. For exposition, suppose that $\mathcal{Y}_{\text{obs}}$ is a measurable vector space and consider the common additive-noise model

$$Y = \mathcal{G}(\theta) + \eta, \tag{15}$$

where, on an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$, $Y, \eta : \Omega \to \mathcal{Y}_{\text{obs}}$ are measurable random variables and η is independent of θ. The certification framework does not rely on this additive form: it only requires a conditional observation law $P(dy \mid \theta)$. Non-additive and implicit observation models, together with the corresponding measurability conditions, are described in Appendix E.

After training on N simulations and observing $Y = y$, an SBI estimator produces an approximate posterior $Q_{N,y} \in \mathcal{P}(\Theta)$. We then condition this learned posterior on the certified set Γ defined in Definition 1 $Q^{\Gamma}_{N,y} := Q_{N,y}(\cdot \mid \Gamma)$. This certification step is applied after posterior training and therefore requires no retraining of the SBI estimator.

The *retained posterior mass* is $a_{N,y} := Q_{N,y}(\Gamma)$, i.e., the probability that a learned-posterior draw passes the certificate. Algorithm 1 estimates this probability from an independent pilot sample and then generates certified draws by rejection sampling. Since each proposal is accepted with probability $a_{N,y}$, obtaining one certified draw requires $1/a_{N,y}$ proposals on average. Theorem 5 further gives $a_{N,y} \geq \alpha - r_N$, so the same retained mass that controls the statistical effect of conditioning also determines its expected sampling cost. A *certificate margin* records the model-specific slack in the sufficient condition used for certification.

5 Experiments

The experiments are organized around the three guarantees in Section 3: coverage of the certified posterior, accuracy of approximate states, and computational cost. We use *posterior mass* in its

Algorithm 1 Certified SBI

Require: Learned posterior $Q_{N,y}$ for data y; certificate C; pilot size K_0; target number M of certified draws.

1: Initialize certification using the fixed learned posterior $Q_{N,y}$.
2: Draw $\widetilde{\theta}^{(1:K_0)} \overset{\text{iid}}{\sim} Q_{N,y}$, evaluate C and margins, and set $\widehat{a} = K_0^{-1} \sum_j C(\widetilde{\theta}^{(j)})$.
3: Set $m \leftarrow 0$, $k \leftarrow 0$.
4: **while** $m < M$ **do**
5: Set $k \leftarrow k + 1$; draw $\theta^{(k)} \sim Q_{N,y}$ and evaluate C.
6: **if** $C(\theta^{(k)}) = 1$ **then**
7: Store the draw and margin; $m \leftarrow m + 1$.
8: **end if**
9: **end while**
10: Return accepted draws, $\widehat{a}$, margins, k, and runtime.

standard probabilistic sense: for a measurable set A, posterior mass is the probability assigned to A by the posterior. Thus retained posterior mass is $Q_{N,y}(\Gamma)$, not a physical mass of the PDE state.

We test posterior mass retained by certification, posterior shift, residual effectivity under solver refinement, and the coverage–cost trade-off as the certificate is tightened. We report retained mass, TV, posterior-mean shift, proposals per accepted sample, and certificate runtime normalized by simulator cost. Ordinary SBI has zero post-training validity-check overhead but no corresponding theorem-backed guarantee. Full protocols are in Appendix I. The certification interface is not tied to the solver used in this benchmark: it requires a bound on the continuous defect rather than a particular numerical architecture. Appendix I analyzes this interface for additional scientific models and details how mesh solvers, PINNs, and neural operators can provide the residual quantities required by the certificates.

5.1 Controlled Nonlinear PDE Benchmark

We consider a Route I benchmark based on the periodic semilinear heat equation

$$\partial_t u - 0.08\, \partial_x^2 u = \lambda u^2, \qquad u(0,x) = A(1 + 0.35\cos(2\pi x)), \tag{16}$$

where $u = u(t,x)$ is the time-dependent state on the one-dimensional periodic domain and $\theta = (A, \lambda)$ is the unknown parameter. We generate synthetic data from the ground-truth parameter $\theta_\star = (0.75, 1.20)$. For three prescribed space–time observation locations (t_j, x_j), $j = 1, 2, 3$, define

$$\mathcal{Y}_{\text{obs}} = \mathbb{R}^3, \qquad \mathcal{G}(\theta) = \big(u^\theta(t_1,x_1), u^\theta(t_2,x_2), u^\theta(t_3,x_3)\big),$$

where u^θ is evaluated using a reference discretization. The observed data are generated according to

$$y = \mathcal{G}(\theta_\star) + \eta, \qquad \eta \sim \mathcal{N}(0, 0.025^2 I_3).$$

Thus the three observation errors are independent Gaussian variables with standard deviation 0.025.

The prior π in Algorithm 1 is uniform on $\Theta = [0.35, 1.15] \times [0.2, 2.2]$. We evaluate the posterior on a transparent 21×21 parameter grid so that the effect of certification is not confounded with neural density-estimation error. On $\mathcal{X} = C([0,T]; L^\infty)$ with $T = 0.15$, heat-semigroup contractivity yields the Route I sufficient condition

$$4\lambda T A(1 + 0.35) < 1. \tag{17}$$

Hence every parameter satisfying (17) is certified by Theorem 1, which guarantees a unique solution in the corresponding certified solution class. Appendix I derives (17) directly from the Route I self-map and contraction conditions.

At the ground-truth parameter $\theta_\star = (0.75, 1.20)$, the certificate retains $a_{N,y} = 71.1\%$ of the posterior mass. Conditioning changes the posterior mean from $(0.758, 1.392)$ to $(0.769, 1.177)$. Since $\mathrm{TV}(Q_{N,y}^\Gamma, Q_{N,y}) = 1 - a_{N,y}$ for conditioning on Γ, the resulting posterior shift is 0.289. The same retained mass implies an average rejection-sampling cost of $1/a_{N,y} = 1.41$ proposals per certified draw.

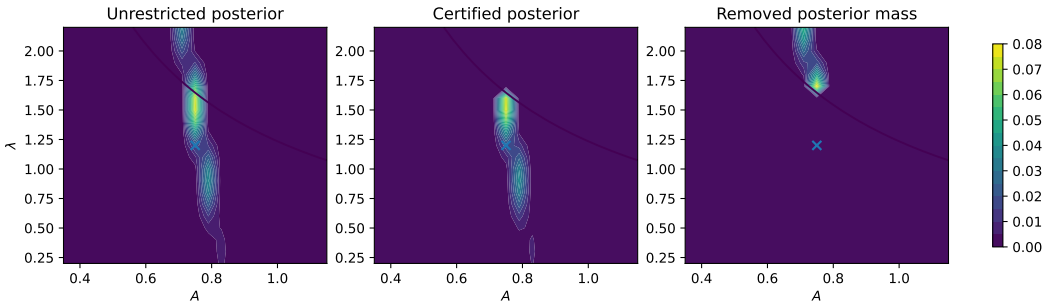


Figure 2: **Certified SBI separates statistical plausibility from mathematical coverage.** Left: the unrestricted posterior $Q_{N,y}$ inferred from the synthetic observations. Center: the certified posterior $Q^{\Gamma}_{N,y} = Q_{N,y}(\cdot \mid \Gamma)$ obtained by retaining only parameters satisfying the Route I certificate (17). Right: posterior mass that is statistically plausible under $Q_{N,y}$ but not covered by this sufficient certificate. The curve marks the boundary $4\lambda TA(1 + 0.35) = 1$.

Certification-threshold ablation. We vary the acceptance threshold from $4\lambda TA(1 + 0.35) < 1$ to the stricter criteria $4\lambda TA(1 + 0.35) < \tau$ with $\tau = 0.85$ and 0.70. The retained posterior mass decreases from 0.711 to 0.405 and 0.266, while the expected proposals per certified draw increase from 1.41 to 2.47 and 3.75, respectively. This ablation makes the coverage–computation trade-off explicit: requiring parameters to lie farther inside the certified regime reduces posterior coverage and increases rejection cost.

6 Related Work

Simulation-based inference. Modern SBI replaces intractable likelihood evaluation by simulation and learned conditional densities or density ratios. Neural posterior estimation, neural likelihood estimation, and neural ratio estimation provide complementary approaches, while benchmarking, sequential, and multifidelity methods improve diagnostics or simulation efficiency (Papamakarios et al., 2019; Greenberg et al., 2019; Hermans et al., 2020; Cranmer et al., 2020; Lueckmann et al., 2021; Krouglova et al., 2026). Our certification layer addresses a complementary question: after an SBI estimator has been trained, is an individual posterior draw covered by a mathematical validity theorem for the underlying continuous scientific model?

Calibration, robustness, and likelihood-free acceptance. Simulation-based calibration evaluates repeated-sampling correctness, posterior predictive checks compare observed and replicated data, and robust SBI addresses model misspecification (Talts et al., 2018; Gelman et al., 1996; Huang et al., 2023). Approximate Bayesian computation can also accept or weight parameters according to discrepancies between simulated and observed data (Beaumont et al., 2002; Sisson et al., 2018). Certified SBI uses a different acceptance criterion: membership in Γ is determined by sufficient conditions from a mathematical validity theorem for the state equation, rather than by agreement between simulated and observed data. The retained mass $Q_{N,y}(\Gamma)$ therefore measures the learned-posterior probability covered by the selected analytical guarantee.

Scientific machine learning and learned PDE solvers. PINNs, DeepONets, Fourier neural operators, and related neural operators learn states or solution maps for differential equations (Raissi et al., 2019; Lu et al., 2021; Li et al., 2021; Kovachki et al., 2023; Li et al., 2024). Their training or test losses are not, in general, the continuous residual bounds required by our certification theorems. Certified SBI separates these layers: a learned solver can be verified when its architecture-specific approximation information provides a rigorous bound on the continuous defect required by Theorem 3 or Theorem 4.

Bayesian inverse problems and verified numerics. Bayesian inverse-problem theory establishes existence, stability, and approximation properties of posterior measures under conditions on the parameter-to-observation map (Stuart, 2010; Dashti and Stuart, 2017). A posteriori PDE analysis and

verified numerics instead control approximate states through residual estimates, stability bounds, and computer-assisted verification techniques (Repin, 2008; Nakao et al., 2019). Certified SBI connects these levels at the posterior-sample scale: deterministic state-space certificates define a measurable conditioning event, while our statistical results quantify how conditioning affects posterior accuracy, downstream quantities, and rejection cost.

7 Discussion and Limitations

Certified SBI provides sufficient rather than necessary guarantees. A rejected parameter may still define a well-posed scientific model; rejection only means that the selected certificate does not establish the required guarantee. Consequently, a conservative analytical theorem can exclude substantial posterior mass even when the underlying model is valid. This limitation is visible through the retained mass $Q_{N,y}(\Gamma)$, which should therefore be reported together with certified posterior summaries.

Theorem-backed numerical certification also requires verified bounds on the continuous defect and on discretization, quadrature, and roundoff errors. Obtaining these bounds can be more expensive than evaluating an ordinary solver diagnostic, and low retained mass further increases rejection-sampling cost. The method therefore trades posterior coverage and additional computation for a theorem-backed guarantee on retained draws.

Finally, certification changes the inferential target from the unrestricted posterior to its conditioning on Γ. The resulting shift is quantified by our posterior-stability and end-to-end bounds, but it cannot be eliminated without enlarging the certified region. Importantly, this also separates improvements in analysis from improvements in inference: a sharper well-posedness or stability theorem can enlarge Γ and recover posterior coverage without retraining the SBI estimator. Further extensions and multi-certificate strategies are discussed in Appendix J.

8 Conclusion

We introduced certified SBI, a modular layer that conditions a learned posterior on parameters covered by deterministic validity certificates for nonlinear scientific models. The framework separates three issues that posterior accuracy alone does not distinguish: approximation of the target posterior, mathematical coverage of the continuous state equation, and numerical error in the state used downstream. The resulting guarantees provide sample-wise well-posedness and solver-error control alongside bounds for conditioned-posterior accuracy, posterior summaries, decisions, and rejection-sampling cost.

The controlled nonlinear PDE experiment illustrates the central validity–coverage–cost trade-off. Posterior mass can remain statistically plausible while lying outside the region covered by the selected mathematical theorem; certification makes this distinction explicit. The retained posterior mass quantifies this coverage and simultaneously determines the expected rejection cost, while tightening the certificate exposes the price of requiring parameters to lie farther inside the certified regime. Additional appendix results examine Route II residual verification and detail the solver-specific steps required for classical and learned PDE solvers.

Certified SBI is deliberately modular: improving the SBI estimator can reduce posterior approximation error, reducing the certified solver defect tightens the corresponding state-error bound, and improving the underlying mathematical theorem can enlarge the certified set without retraining the inference model. The framework makes these distinct limitations measurable through certified coverage, posterior distortion, and verification cost. This separation provides a practical route for combining learned Bayesian inference with explicit mathematical guarantees in nonlinear scientific models.

AI USE STATEMENT

We used generative AI tools to assist with writing and copyediting the manuscript, searching for and identifying related literature, and drafting sections of the paper. We also used these tools to assist with proof construction for technical lemmas and to check completed theorems and lemmas for correctness. We have reviewed all AI-assisted work. Every AI-assisted proof step was independently verified by the author against the underlying mathematics, and AI-assisted drafts and text edits were checked for accuracy before inclusion. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REFERENCES

H. Bahouri, J.-Y. Chemin, and R. Danchin. *Fourier Analysis and Nonlinear Partial Differential Equations*. Springer, 2011.

D. Huang, A. Bharti, A. H. Souza, L. Acerbi, and S. Kaski. Learning robust statistics for simulation-based inference under model misspecification. In *Advances in Neural Information Processing Systems*, 2023.

A. N. Krouglova, H. Johnson, B. Confavreux, M. Deistler, and P. J. Goncalves. Multifidelity simulation-based inference for computationally expensive simulators. In *International Conference on Learning Representations*, 2026.

M. A. Beaumont, W. Zhang, and D. J. Balding. Approximate Bayesian computation in population genetics. *Genetics*, 162(4):2025–2035, 2002.

A. Gelman, X.-L. Meng, and H. Stern. Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica*, 6(4):733–760, 1996.

S. A. Sisson, Y. Fan, and M. A. Beaumont, editors. *Handbook of Approximate Bayesian Computation*. Chapman and Hall/CRC, 2018.

S. Talts, M. Betancourt, D. Simpson, A. Vehtari, and A. Gelman. Validating Bayesian inference algorithms with simulation-based calibration. *arXiv:1804.06788*, 2018.

T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley, second edition, 2006.

S. Banach. Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales. *Fundamenta Mathematicae*, 3:133–181, 1922.

H. Brezis. *Functional Analysis, Sobolev Spaces and Partial Differential Equations*. Springer, 2011.

F. E. Browder. Variational boundary value problems for quasi-linear elliptic equations. *Proceedings of the National Academy of Sciences*, 50:31–37, 1963.

K. Cranmer, J. Brehmer, and G. Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, 2020.

M. Dashti and A. M. Stuart. The Bayesian approach to inverse problems. In *Handbook of Uncertainty Quantification*. Springer, 2017.

D. S. Greenberg, M. Nonnenmacher, and J. H. Macke. Automatic posterior transformation for likelihood-free inference. In *Proceedings of ICML*, 2019.

D. Henry. *Geometric Theory of Semilinear Parabolic Equations*. Springer, 1981.

J. Hermans, V. Begy, and G. Louppe. Likelihood-free MCMC with amortized approximate ratio estimators. In *Proceedings of ICML*, 2020.

N. Kovachki, Z. Li, B. Liu, K. Azizzadenesheli, K. Bhattacharya, A. Stuart, and A. Anandkumar. Neural operator: Learning maps between function spaces with applications to PDEs. *Journal of Machine Learning Research*, 24(89):1–97, 2023.

Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021.

Z. Li, H. Zheng, N. Kovachki, D. Jin, H. Chen, B. Liu, K. Azizzadenesheli, and A. Anandkumar. Physics-informed neural operator for learning partial differential equations. *ACM/IMS Journal of Data Science*, 1(3), 2024.

P. Lindqvist. Notes on the p-Laplace equation. Report 161, University of Jyväskylä, 2017.

L. Lu, P. Jin, G. Pang, Z. Zhang, and G. E. Karniadakis. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3:218–229, 2021.

J.-M. Lueckmann, D. Boelts, D. S. Greenberg, P. J. Goncalves, and J. H. Macke. Benchmarking simulation-based inference. In *Proceedings of AISTATS*, 2021.

G. J. Minty. Monotone (nonlinear) operators in Hilbert space. *Duke Mathematical Journal*, 29:341–346, 1962.

M. T. Nakao, M. Plum, and Y. Watanabe. *Numerical Verification Methods and Computer-Assisted Proofs for Partial Differential Equations*. Springer, 2019.

G. Papamakarios, D. Sterratt, and I. Murray. Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows. In *Proceedings of AISTATS*, 2019.

A. Pazy. *Semigroups of Linear Operators and Applications to Partial Differential Equations*. Springer, 1983.

M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.

S. Repin. *A Posteriori Estimates for Partial Differential Equations*. De Gruyter, 2008.

R. E. Showalter. *Monotone Operators in Banach Space and Nonlinear Partial Differential Equations*. American Mathematical Society, 1997.

A. M. Stuart. Inverse problems: A Bayesian perspective. *Acta Numerica*, 19:451–559, 2010.

E. Zeidler. *Nonlinear Functional Analysis and Its Applications II/A: Linear Monotone Operators*. Springer, 1990.

M. D. Reid and R. C. Williamson. Information, divergence and risk for binary experiments. *Journal of Machine Learning Research*, 12:731–817, 2011.

G. L. Gilardoni. On Pinsker's and Vajda's type inequalities for Csiszár's f-divergences. *IEEE Transactions on Information Theory*, 56(11):5377–5386, 2010.

I. Sason and S. Verdú. f-divergence inequalities. *IEEE Transactions on Information Theory*, 62(11):5973–6006, 2016.

A. Müller. Integral probability metrics and their generating classes of functions. *Advances in Applied Probability*, 29(2):429–443, 1997.

B. K. Sriperumbudur, K. Fukumizu, A. Gretton, B. Schölkopf, and G. R. G. Lanckriet. On the empirical estimation of integral probability metrics. *Electronic Journal of Statistics*, 6:1550–1599, 2012.

R. Agrawal and T. Horel. Optimal bounds between f-divergences and integral probability metrics. *Journal of Machine Learning Research*, 22(128):1–59, 2021.

F. Liese and I. Vajda. On divergences and informations in statistics and information theory. *IEEE Transactions on Information Theory*, 52(10):4394–4412, 2006.

A. Hauptmann, V. Kolehmainen, N. M. Mach, T. Savolainen, A. Seppänen, and S. Siltanen. Open 2D Electrical Impedance Tomography data archive. *arXiv preprint arXiv:1704.01178*, 2017. Dataset DOI: 10.5281/zenodo.1203914.

W. E and B. Yu. The Deep Ritz method: A deep learning-based numerical algorithm for solving variational problems. *Communications in Mathematics and Statistics*, 6(1):1–12, 2018.

A Notation and Mathematical Setting

Table 1 collects the symbols used throughout the paper. The table is intentionally explicit because the framework combines notation from SBI, functional analysis, and PDE theory.

Table 1: Notation used in the paper.

Symbol	Meaning
Θ	Parameter space with σ-algebra $\mathcal{B}(\Theta)$.
$\mathcal{P}(\Theta)$	Probability measures on $(\Theta, \mathcal{B}(\Theta))$.
θ	Unknown model parameter.
π	Prior distribution for θ.
$(\mathcal{Y}_{\mathrm{obs}}, \Sigma_{\mathrm{obs}})$	Measurable observation space and its σ-algebra.
Y, y	Random observation Y and its realized value $y \in \mathcal{Y}_{\mathrm{obs}}$.
$\mathcal{G}$	Measurable parameter-to-observation map $\Theta \to \mathcal{Y}_{\mathrm{obs}}$.
$P(dy \mid \theta)$	Conditional observation law (simulator kernel) at parameter θ.
η	Observation noise, independent of θ in the additive model (15).
$\mathcal{U}$	Banach data space (initial, boundary, or forcing data).
u_0	Initial datum for an evolution problem.
$\mathcal{X}$	Banach solution space with norm $\|\cdot\|_{\mathcal{X}}$.
$\mathcal{X}^*$	Continuous dual of $\mathcal{X}$ with norm $\|\cdot\|_{\mathcal{X}^*}$.
$\langle \ell, v \rangle$	Dual pairing between $\ell \in \mathcal{X}^*$ and $v \in \mathcal{X}$.
u^θ	Exact state associated with parameter θ.
T_θ	Fixed-point solution map $\mathcal{X} \to \mathcal{X}$.
$\mathcal{R}_\theta$	Continuous residual operator, typically $\mathcal{X} \to \mathcal{X}^*$ or $\mathcal{X} \to \mathcal{Y}$.
$\mathcal{E}_\theta$	Fixed-point defect operator $\mathcal{E}_\theta(v) = v - T_\theta(v) \in \mathcal{X}$.
r^θ	Route II residual $r^\theta := \mathcal{R}_\theta(v^\theta) = \mathcal{A}_\theta(v^\theta) - f_\theta \in \mathcal{X}^*$.
$\overline{B}_{\mathcal{X}}(0, R_\theta)$	Certified closed ball in $\mathcal{X}$.
$C(\theta)$	Certificate indicator: 1 if the chosen theorem verifies θ, otherwise 0.
Γ	Certified parameter set $\{\theta : C(\theta) = 1\}$.
$Q_{N,y}$	Learned posterior approximation after N simulations and observation y.
$Q^\Gamma_{N,y}$	Certified learned posterior $Q_{N,y}(\cdot \mid \Gamma)$.
Π_y	Reference posterior used in the stability theorem.
Π^Γ_y	Certified reference posterior $\Pi_y(\cdot \mid \Gamma)$.
$\mathrm{Cert}_\Gamma(Q)$	Conditional distribution $Q(\cdot \mid \Gamma)$.
$a_{N,y}$	Retained posterior mass $Q_{N,y}(\Gamma)$; acceptance probability for independent rejection sampling.
$\widehat{a}$	Pilot Monte Carlo estimate of $a_{N,y}$.
K_0, M	Pilot size and requested number of accepted draws in Algorithm 1.
TV	Total-variation metric on $\mathcal{P}(\Theta)$.
D	Divergence $\mathcal{P}(\Theta)^2 \to [0, \infty]$ used to quantify posterior discrepancy.
ψ_D	Calibration function satisfying $\mathrm{TV}(P, Q) \le \psi_D(D(Q\|P))$.
KL	Kullback–Leibler divergence, a particular choice of D.
α	Lower bound on reference certified mass $\Pi_y(\Gamma)$.
r_N	TV calibration radius $\psi_D(\varepsilon_N)$ in Theorem 5.
λ_θ	Contraction factor on the certified state region.
ω_θ	Stability modulus for the residual route.
m_θ	Positive quantitative stability constant in $\omega_\theta(s) \ge m_\theta s^p$.

Banach dual. For completeness, we use the standard Banach-space duality conventions of Brezis (2011); Zeidler (1990):

$$\mathcal{X}^* = \{\ell : \mathcal{X} \to \mathbb{R} : \ell \text{ is linear and continuous}\}, \qquad \|\ell\|_{\mathcal{X}^*} = \sup_{\|v\|_{\mathcal{X}} \le 1} |\ell(v)| .$$

When $\mathcal{R}_\theta : \mathcal{X} \to \mathcal{X}^*$, the duality inequality $|\langle \ell, v \rangle| \le \|\ell\|_{\mathcal{X}^*} \|v\|_{\mathcal{X}}$ is the basic step that turns a residual bound into a state-error bound.

A.1 Well-posedness and stability terminology

Definition 3 (Well-posed model instance)**.** *Let $(\mathcal{X}, \|\cdot\|_{\mathcal{X}})$ be a Banach space in which the state equation is posed, and let $\mathcal{S}_\theta \subseteq \mathcal{X}$ denote the solution class associated with parameter θ. We say*

that the model is well posed in $\mathcal{S}_\theta$ *if a solution exists, is unique in* $\mathcal{S}_\theta$*, and depends continuously on the model data with respect to the topology specified by the certificate.*

In Route I the certified solution class is the closed ball $\overline{B}_{\mathcal{X}}(0, R_\theta)$ on which the self-map and contraction inequalities hold. In Route II stability is quantified by the modulus ω_θ in (5); positivity away from zero implies uniqueness once existence is known.

B Meaning of the Model Equations

This section makes precise the two equations used in the main text.

B.1 Parameter-dependent fixed-point map

Let $(\Theta, \mathcal{B}(\Theta))$ be the measurable parameter space used for Bayesian inference, and let $(\mathcal{X}, \|\cdot\|_{\mathcal{X}})$ be a Banach space in which the scientific state is sought. A parameter-dependent fixed-point model is a map

$$T : \Theta \times \mathcal{X} \to \mathcal{X}, \qquad (\theta, u) \mapsto T(\theta, u),$$

and, after fixing $\theta \in \Theta$, we write $T_\theta(u) := T(\theta, u)$. The parameter θ may encode physical coefficients, initial or boundary data, forcing parameters, or other unknowns inferred by SBI. For posterior certification, T need not be globally contractive on $\mathcal{X}$; Route I only requires the parameter-specific self-map and contraction conditions on the candidate region $\overline{B}_{\mathcal{X}}(0, R_\theta)$.

Fixed-point equality. For a fixed parameter θ, a state $u \in \mathcal{X}$ solves the fixed-point model when $T_\theta(u)$ is well defined in $\mathcal{X}$ and

$$u = T_\theta(u) \quad \text{in } \mathcal{X},$$

or equivalently,

$$\|u - T_\theta(u)\|_{\mathcal{X}} = 0.$$

Thus the equation is interpreted in the topology of the state space $\mathcal{X}$, rather than as a pointwise identity unless the chosen deterministic model provides additional regularity.

A principal example is a semilinear evolution equation

$$\partial_t u - A_\theta u = F_\theta(u), \qquad u(0) = u_0.$$

Assume that the linear operator A_θ generates an evolution family $S_\theta(t)$ for which the variation-of-constants formula is valid (Pazy, 1983; Henry, 1981). This formulation is more general than a diffusion equation: when A_θ is Laplace-type, $S_\theta(t)$ describes diffusive evolution, while other choices of A_θ generate different linear dynamics. The corresponding fixed-point map is

$$T_\theta(u)(t) = S_\theta(t)u_0 + \int_0^t S_\theta(t-s)F_\theta(u(s))\,ds = L_\theta(u_0)(t) + \mathcal{N}_\theta(u)(t). \tag{18}$$

Here $L_\theta(u_0) \in \mathcal{X}$ is the contribution of the linear evolution and $\mathcal{N}_\theta(u) \in \mathcal{X}$ is the nonlinear contribution. In time-dependent applications, $\mathcal{X}$ is chosen as a time–space Banach space appropriate to the deterministic well-posedness theorem; pointwise differentiability of u is not required unless that theorem supplies it.

Residual equality. When $\mathcal{R}_\theta : \mathcal{X} \to \mathcal{X}^*$, the residual equation (2) means

$$\mathcal{R}_\theta(u) = 0 \qquad \text{in } \mathcal{X}^*.$$

Equivalently,

$$\langle \mathcal{R}_\theta(u), \varphi \rangle_{\mathcal{X}^*, \mathcal{X}} = 0 \qquad \text{for every } \varphi \in \mathcal{X}.$$

This weak formulation allows Route II to certify steady-state models whose solutions need not possess pointwise derivatives, provided the residual is well defined as an element of $\mathcal{X}^*$.

C Function Spaces Used by the Certificates

A function space is needed to specify the admissible scientific states, the notion of distance between two states, and the norm in which stability and residual-to-error guarantees are stated. For each model, we therefore choose a Banach space $(\mathcal{X}, \|\cdot\|_{\mathcal{X}})$ in which the deterministic existence, uniqueness, and stability theorem underlying the certificate is valid. The appropriate choice depends on the scientific model and on the available theorem. For example, a time-dependent diffusion problem may use a space controlling the solution uniformly over time and space, whereas a steady variational problem naturally uses a Sobolev space controlling spatial derivatives. Once this deterministic state space is fixed, the statistical conditioning step of certified SBI is unchanged. We summarize below the function spaces used in this paper.

Let $\Omega \subseteq \mathbb{R}^d$ be a spatial domain. For $1 \le q < \infty$,

$$L^q(\Omega) = \left\{ f : \int_\Omega |f(x)|^q \, dx < \infty \right\}, \qquad \|f\|_{L^q} = \left(\int_\Omega |f(x)|^q \, dx \right)^{1/q}.$$

The Sobolev space $W^{k,q}(\Omega)$ additionally requires weak spatial derivatives up to order k to belong to $L^q(\Omega)$, and $H^k(\Omega) = W^{k,2}(\Omega)$ (Brezis, 2011). In particular, $H_0^1(\Omega)$ and $W_0^{1,p}(\Omega)$ are natural state spaces for the steady Route II models considered in this paper: their norms control spatial derivatives, while their dual spaces provide the natural setting for weak PDE residuals.

Besov spaces provide a frequency-resolved alternative to Sobolev spaces and enter this paper in two specific extensions. First, some scale-invariant well-posedness theorems for incompressible flow formulate their sufficient conditions in Besov norms, so such a theorem can induce a Besov-based certificate. Second, Appendix E considers integral probability metrics whose test functions may be chosen from a Besov class. A representative homogeneous Besov norm is

$$\|f\|_{\dot{B}^s_{p,q}} = \left(\sum_j 2^{jsq} \, \|\Delta_j f\|_{L^p}^q \right)^{1/q},$$

with the usual modification for $q = \infty$, where Δ_j extracts the jth dyadic frequency band (Bahouri et al., 2011). Thus Besov spaces are not an additional requirement of the certification framework; they provide one possible certificate or test-function geometry when required by the corresponding deterministic or statistical result.

For the controlled time-dependent benchmark in Section 5, we use

$$\mathcal{X} = C([0,T]; L^\infty(\Omega)),$$

where

$$C([0,T]; L^\infty(\Omega)) = \{u : [0,T] \to L^\infty(\Omega) : u \text{ is continuous with respect to the } L^\infty \text{ norm}\},$$

equipped with

$$\|u\|_{\mathcal{X}} = \|u\|_{C_t L_x^\infty} := \sup_{0 \le t \le T} \|u(t,\cdot)\|_{L^\infty(\Omega)}.$$

This choice is not arbitrary: heat-flow contractivity controls the linear evolution in this norm, while the quadratic nonlinearity is locally Lipschitz on bounded subsets of $\mathcal{X}$. These two properties yield the explicit Route I self-map and contraction certificate derived in Appendix I.

D Mathematical Foundations of Certification

This appendix makes the deterministic certificate oracles explicit. The main text treats them through a common interface; here we state why the two routes are needed and what each guarantees.

D.1 Well-Posedness, Stability, and Banach's Theorem

Definition 3 is relative to a solution class because PDE uniqueness can hold in one topology while remaining unknown in another. For this paper, a certified problem is *stable* on a solution class $\mathcal{S}_\theta$ if

state perturbations are controlled quantitatively by either a contraction margin as in (4) or a residual modulus as in (5). For perturbations of external data, stability means continuity of the solution map in the data topology specified by the deterministic theorem. Thus "stable" is always tied to an explicit norm and modulus, never to an informal numerical notion.

Theorem 7 (Banach fixed-point theorem (Banach, 1922)). *Let (S, d) be a nonempty complete metric space and let $T : S \to S$ satisfy $d(Tu, Tv) \le \lambda d(u, v)$ for all $u, v \in S$ and some $0 \le \lambda < 1$. Then T has a unique fixed point $u^* \in S$, and the iterates $u_{n+1} = T(u_n)$ converge to u^* for every $u_0 \in S$.*

For Theorem 1, take $S = \overline{B}_{\mathcal{X}}(0, R_\theta)$. The self-map inequality gives $T_\theta(S) \subseteq S$ and the Lipschitz inequality gives a strict contraction. Since a closed ball in a Banach space is complete, the theorem applies. This proves existence and uniqueness *inside the certified state region*; no uniqueness claim is made in a larger weak-solution class.

D.2 Definitions for the residual well-posedness theorem

Let $\mathcal{X}$ be a reflexive Banach space and $\mathcal{A}_\theta : \mathcal{X} \to \mathcal{X}^*$. The assumptions in Theorem 2 mean:

- $\mathcal{A}_\theta$ is *bounded on bounded sets* if for every $R > 0$ there exists $C_R < \infty$ such that $\|\mathcal{A}_\theta(u)\|_{\mathcal{X}^*} \le C_R$ whenever $\|u\|_{\mathcal{X}} \le R$;
- it is *hemicontinuous* if for all $u, v, w \in \mathcal{X}$, the scalar map $t \mapsto \langle \mathcal{A}_\theta(u + tv), w\rangle_{\mathcal{X}^*,\mathcal{X}}$ is continuous;
- it is *coercive* if
$$\frac{\langle \mathcal{A}_\theta(u), u\rangle_{\mathcal{X}^*,\mathcal{X}}}{\|u\|_{\mathcal{X}}} \to +\infty \quad \text{as } \|u\|_{\mathcal{X}} \to \infty;$$
- it is *uniformly monotone with modulus ω_θ* if
$$\langle \mathcal{A}_\theta(u) - \mathcal{A}_\theta(v), u - v\rangle_{\mathcal{X}^*,\mathcal{X}} \ge \omega_\theta(\|u - v\|_{\mathcal{X}})$$
for all $u, v \in \mathcal{X}$, where $\omega_\theta(s) > 0$ for $s > 0$. This implies strict monotonicity and is exactly the quantitative stability condition used in Theorem 2.

These are deterministic properties of the scientific operator; they are not learned from data.

Theorem 8 (Browder–Minty existence theorem (Browder, 1963; Minty, 1962; Showalter, 1997)). *Let $\mathcal{X}$ be reflexive and let $A : \mathcal{X} \to \mathcal{X}^*$ be bounded on bounded sets, hemicontinuous, coercive, and monotone. Then A is surjective onto $\mathcal{X}^*$. If A is strictly monotone, the solution of $A(u) = f$ is unique for every $f \in \mathcal{X}^*$.*

This theorem is included because it is the deterministic existence result used to instantiate Route II. The stability condition (5) is exactly uniform monotonicity of $\mathcal{A}_\theta$ because the source term f_θ cancels in residual differences. Positivity of the modulus implies strict monotonicity and therefore uniqueness.

D.3 Concrete Route II Operators and ML-Relevant Examples

The decomposition $\mathcal{R}_\theta(u) = \mathcal{A}_\theta(u) - f_\theta$ is useful whenever the scientific state is characterized by a steady balance, variational optimum, or weak conservation law. These models occur in inverse problems, fluid mechanics, and neural variational solvers.

Nonlinear diffusion and conductivity. For $\mathcal{X} = W_0^{1,p}(\Omega)$, $1 < p < \infty$,

$$\langle \mathcal{A}_\theta(u), v\rangle_{\mathcal{X}^*,\mathcal{X}} = \int_\Omega a_\theta(x)|\nabla u|^{p-2}\nabla u \cdot \nabla v \, dx, \qquad \langle f_\theta, v\rangle_{\mathcal{X}^*,\mathcal{X}} = \int_\Omega g_\theta v \, dx.$$

Uniform positivity of a_θ gives a monotone diffusion operator. For $p = 2$ this includes conductivity/Poisson inverse problems; for $p \ne 2$ it gives nonlinear diffusion. Here "diffusion" refers to physical transport in the PDE and is distinct from generative diffusion models in machine learning (Lindqvist, 2017; Showalter, 1997).

Reaction–diffusion equilibria. A steady reaction–diffusion model can be written weakly as

$$\langle \mathcal{A}_\theta(u), v\rangle_{\mathcal{X}^*,\mathcal{X}} = \int_\Omega a_\theta(x,u,\nabla u)\cdot\nabla v + b_\theta(x,u)v\,dx.$$

When the diffusion and reaction terms satisfy quantitative monotonicity, Route II turns the weak residual of an approximate equilibrium into a state error bound. The time-dependent version may instead be treated by Route I; the two routes correspond to different deterministic descriptions of the same physical mechanism.

Steady viscous flow. On the divergence-free space

$$\mathcal{X}_{\mathrm{div}} = \{u \in W_0^{1,p}(\Omega)^d : \nabla\cdot u = 0\},$$

a representative generalized-Newtonian viscous operator is

$$\langle \mathcal{A}_\theta(u), v\rangle_{\mathcal{X}^*_{\mathrm{div}},\mathcal{X}_{\mathrm{div}}} = \int_\Omega 2\mu_\theta(x)|Du|^{p-2}Du : Dv\,dx, \qquad Du = \tfrac{1}{2}(\nabla u + \nabla u^\top).$$

Under the usual positivity and monotonicity conditions on the viscosity, this has the Route II structure. It complements the time-dependent incompressible-flow Route I example: the route is determined by the available stability theorem, not by the application label "fluid."

Variational neural solvers. Many steady PDEs minimize an energy $J_\theta : \mathcal{X} \to \mathbb{R}$. If J_θ is differentiable and strongly convex, then $\mathcal{A}_\theta = J'_\theta$ is strongly monotone and $J'_\theta(u) = 0$ is a Route II residual equation. A neural variational method such as the Deep Ritz method (E and Yu, 2018) can supply an approximate state v^θ, while certified SBI evaluates the continuous dual residual $J'_\theta(v^\theta)$ rather than treating the training loss itself as a state-error certificate. Thus the network proposes a state and the deterministic stability inequality converts its continuous residual into a theorem-backed error bar.

D.4 Why Route II Complements Fixed-Point Certification

Route I is effective when the scientific model admits a useful contraction map, but many inverse problems are posed more naturally through a residual or variational equation. Route II is included so that certification does not depend on rewriting every simulator as a time-evolution fixed-point problem. For steady conductivity, nonlinear diffusion, and related equilibrium models, the directly computable object is typically $\mathcal{R}_\theta(u) = 0$ together with a stability inequality. This formulation also matches the residual information produced by common variational discretizations.

One implementation uses monotone-operator theory. Let $\mathcal{X}$ be reflexive and $\mathcal{A}_\theta : \mathcal{X} \to \mathcal{X}^*$ be bounded on bounded sets, hemicontinuous, coercive, and monotone. Browder–Minty theory yields existence under standard hypotheses; strict monotonicity yields uniqueness (Browder, 1963; Minty, 1962; Showalter, 1997). This is only one certificate oracle: inf-sup estimates, Newton–Kantorovich bounds, strong convexity, or verified linearization can replace it without changing the posterior layer.

For the Dirichlet p-Laplacian, $\mathcal{X} = W_0^{1,p}(\Omega)$ and

$$\langle \mathcal{A}(u), v\rangle = \int_\Omega |\nabla u|^{p-2}\nabla u\cdot\nabla v\,dx.$$

For $1 < p < \infty$ this operator is coercive and monotone; standard Dirichlet problems with $f \in W^{-1,p'}(\Omega)$ are therefore natural examples (Lindqvist, 2017). The example demonstrates coverage of stationary nonlinear simulators; it is not a claim that p-Laplace is the primary target application.

D.5 Measurability and Well-Defined Posterior Conditioning

The main text defines the certificate indicator $C : \Theta \to \{0,1\}$ and the certified set

$$\Gamma = \{\theta \in \Theta : C(\theta) = 1\}.$$

Two technical conditions are needed before conditioning is meaningful. First, C must be measurable, so that $\Gamma = C^{-1}(\{1\}) \in \mathcal{B}(\Theta)$. In concrete certificates this follows, for example, when the norms,

stability constants, and verified residual bounds entering the Route I or Route II inequalities depend measurably on θ. This uses the same indicator C as the main text and does not introduce a second certificate score.

Second, the posterior must assign positive mass to Γ. If $Q(\Gamma) = 0$, the conditioning formula has a zero denominator. If $Q \in \mathcal{P}(\Theta)$ and $Q(\Gamma) > 0$, define

$$Q^{\Gamma}(A) = \frac{Q(A \cap \Gamma)}{Q(\Gamma)}, \qquad A \in \mathcal{B}(\Theta).$$

For pairwise disjoint measurable sets $\{A_j\}$,

$$Q^{\Gamma}\left(\bigcup_j A_j\right) = \frac{\sum_j Q(A_j \cap \Gamma)}{Q(\Gamma)} = \sum_j Q^{\Gamma}(A_j),$$

and

$$Q^{\Gamma}(\Theta) = 1, \qquad Q^{\Gamma}(\Gamma) = 1.$$

Hence

$$\mathrm{Cert}_{\Gamma} : \{Q \in \mathcal{P}(\Theta) : Q(\Gamma) > 0\} \to \mathcal{P}(\Theta), \qquad Q \mapsto Q^{\Gamma},$$

is well defined. Applying this argument to $Q_{N,y}$ and Π_y explains the positivity assumptions in Theorem 5 and justifies the statement that a draw from the certified posterior belongs to Γ almost surely.

E Probability Preliminaries

Well-defined observation model. Let $(\Theta, \mathcal{B}(\Theta))$ be the parameter space and let $(\mathcal{Y}_{\mathrm{obs}}, \Sigma_{\mathrm{obs}})$ be a measurable observation space. A measurable parameter-to-observation map $\mathcal{G} : \Theta \to \mathcal{Y}_{\mathrm{obs}}$ associates each parameter with the ideal observation predicted by the continuous scientific model after application of the observation operator. If $\mathcal{Y}_{\mathrm{obs}}$ is a measurable topological vector space, addition is measurable, and $\eta : \Omega_0 \to \mathcal{Y}_{\mathrm{obs}}$ is a measurable noise variable on a probability space $(\Omega_0, \mathcal{F}_0, \mathbb{P}_0)$, then the additive model $Y = \mathcal{G}(\theta) + \eta$ defines an $\mathcal{Y}_{\mathrm{obs}}$-valued random observation. More generally, a Markov kernel $P(\,\cdot \mid \theta)$ from Θ to $(\mathcal{Y}_{\mathrm{obs}}, \Sigma_{\mathrm{obs}})$ covers non-additive, heteroscedastic, discrete, and implicit observation models.

Concrete observation spaces and parameter-to-observation maps. The pair $(\mathcal{Y}_{\mathrm{obs}}, \Sigma_{\mathrm{obs}})$ accommodates several standard observation formats:

- *Finite sensor arrays.* For m scalar measurements, take $\mathcal{Y}_{\mathrm{obs}} = \mathbb{R}^m$ with its Borel σ-algebra. If u^{θ} is the continuous state and $\mathcal{O} : \mathcal{X} \to \mathbb{R}^m$ is a measurable observation operator, then $\mathcal{G}(\theta) = \mathcal{O}(u^{\theta})$.
- *Time series.* For continuously recorded outputs, one may take $\mathcal{Y}_{\mathrm{obs}} = C([0,T]; \mathbb{R}^m)$ and $\mathcal{G}(\theta)(t) = \mathcal{O}(u^{\theta}(t))$. Sampling at $t_1, \ldots, t_J$ instead gives a vector in $\mathbb{R}^{mJ}$.
- *Field-valued data.* For spatial fields or images, one may take $\mathcal{Y}_{\mathrm{obs}} = L^2(D)$ and $\mathcal{G}(\theta) = \mathcal{O}(u^{\theta})$ for a measurable restriction, projection, or imaging operator $\mathcal{O}$.

For the controlled benchmark in Section 5, the observations at finitely many space–time locations are of the first type: $\mathcal{Y}_{\mathrm{obs}} = \mathbb{R}^m$, and $\mathcal{G}(A, \lambda)$ is the vector of noise-free sensor values extracted from the solution of (16). Thus $\mathcal{G}$ combines the parameter-to-state map with the observation operator actually used to generate data.

Divergences on probability measures. A *divergence* in this paper is a map

$$D : \mathcal{P}(\Theta) \times \mathcal{P}(\Theta) \to [0, \infty]$$

satisfying $D(P\|P) = 0$ for every $P \in \mathcal{P}(\Theta)$. This usage follows the standard statistical-information notion of a divergence: it measures discrepancy between probability laws but need not be symmetric or satisfy the triangle inequality; see, e.g., Liese and Vajda (2006) and the classical f-divergence

literature reviewed there. A divergence is *proper* if $D(Q\|P) = 0$ implies $Q = P$. The posterior-stability theorem only uses divergences for which there exists a nondecreasing calibration function $\psi_D : [0, \infty] \to [0, 1]$, with $\psi_D(0) = 0$, such that

$$\mathrm{TV}(P, Q) \le \psi_D(D(Q\|P))$$

for the probability measures under consideration. Thus the abstract quantity D in the main text must quantitatively control total variation through the stated calibration inequality.

Kullback–Leibler and f-divergences. For $Q, P \in \mathcal{P}(\Theta)$ with $Q \ll P$,

$$\mathrm{KL}(Q\|P) = \int_\Theta \log\left(\frac{dQ}{dP}\right) dQ,$$

and $\mathrm{KL}(Q\|P) = +\infty$ otherwise. Thus a finite KL assumption in a theorem implicitly includes absolute continuity. More generally, for a convex $f : [0, \infty) \to \mathbb{R} \cup \{+\infty\}$ with $f(1) = 0$, the Csiszár f-divergence is

$$D_f(Q\|P) = \int f\left(\frac{dQ}{dP}\right) dP$$

when $Q \ll P$, with the standard lower-semicontinuous extension for singular parts. KL is the choice $f(t) = t \log t$. Generalized Pinsker inequalities relate many f-divergences to variational distance (Reid and Williamson, 2011; Gilardoni, 2010; Sason and Verdú, 2016); this is why Theorem 5 is stated for an abstract TV-calibrated divergence rather than KL alone.

Lemma 1 (Pinsker's inequality)**.** *For any probability measures $P, Q \in \mathcal{P}(\Theta)$,*

$$\mathrm{TV}(P, Q) \le \sqrt{\frac{1}{2}\,\mathrm{KL}(Q\|P)}.$$

See, e.g., Cover and Thomas (2006). Hence KL satisfies (11) with $\psi_{\mathrm{KL}}(t) = \sqrt{t/2}$.

Integral probability metrics and Besov test classes. Total variation is an integral probability metric (IPM) over bounded measurable test functions. More generally, for a symmetric function class $\mathcal{F}$, define

$$d_{\mathcal{F}}(P, Q) := \sup_{h \in \mathcal{F}} |\mathbb{E}_P h - \mathbb{E}_Q h|.$$

This viewpoint includes Wasserstein and kernel/MMD discrepancies and provides a natural refinement when downstream observables are smoother than arbitrary bounded functions (Müller, 1997; Sriperumbudur et al., 2012; Agrawal and Horel, 2021). If $\Theta \subset \mathbb{R}^{d_\theta}$ and $\mathcal{F}$ is the unit ball of an inhomogeneous Besov space $B^s_{\infty,1}(\Theta)$, then

$$|\mathbb{E}_P h - \mathbb{E}_Q h| \le \|h\|_{B^s_{\infty,1}}\, d_{B^s_{\infty,1}}(P, Q).$$

Thus Theorem 6 admits a Besov-observable analogue whenever the inference analysis controls this Besov IPM (or a divergence known to dominate it). This is not automatic from the present assumptions, so we do not replace the robust L^∞/TV statement in the main theorem. The advantage is potentially sharper downstream bounds for smooth losses or observables, rather than a stronger claim for arbitrary bounded quantities.

F Complete Proofs of the Main Guarantees

F.1 Proof of Theorem 1

Proof of Theorem 1. Fix $\theta \in \Gamma$ and write $S_\theta := \overline{B}_{\mathcal{X}}(0, R_\theta)$. Since $\mathcal{X}$ is a Banach space and S_θ is closed in $\mathcal{X}$, the metric space $(S_\theta, \|\cdot\|_{\mathcal{X}})$ is complete.

We first show that T_θ maps S_θ into itself. Let $u \in S_\theta$. By the triangle inequality,

$$\begin{aligned}\|T_\theta(u)\|_{\mathcal{X}} &\le \|T_\theta(0)\|_{\mathcal{X}} + \|T_\theta(u) - T_\theta(0)\|_{\mathcal{X}} \\ &\le \|T_\theta(0)\|_{\mathcal{X}} + \sup_{w \in S_\theta} \|T_\theta(w) - T_\theta(0)\|_{\mathcal{X}} \\ &\le R_\theta,\end{aligned}$$

where the final step is (3). Hence $T_\theta(u) \in S_\theta$.

Next, for any $u, v \in S_\theta$, (4) gives

$$\|T_\theta(u) - T_\theta(v)\|_{\mathcal{X}} \le \lambda_\theta \|u - v\|_{\mathcal{X}}, \qquad 0 \le \lambda_\theta < 1.$$

Thus $T_\theta : S_\theta \to S_\theta$ is a strict contraction. By Banach's fixed-point theorem (Theorem 7), there exists a unique $u^\theta \in S_\theta$ satisfying $u^\theta = T_\theta(u^\theta)$.

Now let θ be sampled from $Q^\Gamma_{N,y}$. The notation

$$\theta \sim Q^\Gamma_{N,y}$$

means that θ is a random variable whose probability law is $Q^\Gamma_{N,y}$. By Appendix D.5,

$$\begin{aligned} Q^\Gamma_{N,y}(\Gamma) &= \frac{Q_{N,y}(\Gamma \cap \Gamma)}{Q_{N,y}(\Gamma)} \\ &= \frac{Q_{N,y}(\Gamma)}{Q_{N,y}(\Gamma)} = 1. \end{aligned}$$

Therefore a draw from the certified posterior belongs to Γ with probability one, and the preceding existence-and-uniqueness argument applies to every such draw except on a set of posterior probability zero. □

F.2 Proof of Theorem 2

Proof of Theorem 2. The hypotheses are the standard Browder–Minty assumptions. Because $\mathcal{X}$ is reflexive, $\mathcal{A}_\theta$ is bounded on bounded sets, hemicontinuous, coercive, and monotone, the Browder–Minty theorem implies that $\mathcal{A}_\theta : \mathcal{X} \to \mathcal{X}^*$ is surjective. Hence, for the prescribed $f_\theta \in \mathcal{X}^*$, there exists at least one $u^\theta \in \mathcal{X}$ satisfying $\mathcal{A}_\theta(u^\theta) = f_\theta$.

To prove uniqueness, suppose $u_1, u_2 \in \mathcal{X}$ are two solutions. Then

$$\mathcal{A}_\theta(u_1) - \mathcal{A}_\theta(u_2) = 0.$$

Strict monotonicity gives, whenever $u_1 \neq u_2$,

$$\langle \mathcal{A}_\theta(u_1) - \mathcal{A}_\theta(u_2), u_1 - u_2 \rangle_{\mathcal{X}^*, \mathcal{X}} > 0,$$

whereas the left-hand side is zero because both states solve the same equation. This contradiction implies $u_1 = u_2$. Therefore the solution exists and is unique. □

F.3 Proof of Theorem 3

Proof of Theorem 3. Fix $\theta \in \Gamma$. Let $u = u^\theta$ be the unique solution in $\overline{B}_{\mathcal{X}}(0, R_\theta)$ and let $v = v^\theta \in \overline{B}_{\mathcal{X}}(0, R_\theta)$ be an approximate state. Since $u = T_\theta(u)$ and $\mathcal{E}_\theta(v) = v - T_\theta(v)$,

$$\begin{aligned} u - v &= T_\theta(u) - v \\ &= T_\theta(u) - T_\theta(v) - \mathcal{E}_\theta(v). \end{aligned}$$

Taking the $\mathcal{X}$-norm, applying the triangle inequality, and then using (4),

$$\begin{aligned} \|u - v\|_{\mathcal{X}} &\le \|T_\theta(u) - T_\theta(v)\|_{\mathcal{X}} + \|\mathcal{E}_\theta(v)\|_{\mathcal{X}} \\ &\le \lambda_\theta \|u - v\|_{\mathcal{X}} + \|\mathcal{E}_\theta(v)\|_{\mathcal{X}}. \end{aligned}$$

Because $\lambda_\theta < 1$,

$$(1 - \lambda_\theta) \|u - v\|_{\mathcal{X}} \le \|\mathcal{E}_\theta(v)\|_{\mathcal{X}}.$$

Dividing by $1 - \lambda_\theta > 0$ proves (8). □

Remark 1. *The condition $u^\theta, v^\theta \in \overline{B}_{\mathcal{X}}(0, R_\theta)$ is structural, not merely technical: the contraction estimate is only known on the state region on which the certificate was verified.*

F.4 Proof of Theorem 4

Proof of Theorem 4. Set $u = u^\theta$ and $v = v^\theta$. By (5),

$$\omega_\theta(\|u - v\|_{\mathcal{X}}) \le \langle \mathcal{R}_\theta(u) - \mathcal{R}_\theta(v), u - v \rangle_{\mathcal{X}^*,\mathcal{X}}. \tag{19}$$

Since $\mathcal{R}_\theta(u) = 0$ and $r^\theta = \mathcal{R}_\theta(v)$,

$$\begin{aligned}\langle \mathcal{R}_\theta(u) - \mathcal{R}_\theta(v), u - v \rangle_{\mathcal{X}^*,\mathcal{X}} &= \left\langle -r^\theta, u - v \right\rangle_{\mathcal{X}^*,\mathcal{X}} \\ &= \left\langle r^\theta, v - u \right\rangle_{\mathcal{X}^*,\mathcal{X}}.\end{aligned}$$

The defining inequality for the dual norm yields

$$\left| \left\langle r^\theta, v - u \right\rangle_{\mathcal{X}^*,\mathcal{X}} \right| \le \left\| r^\theta \right\|_{\mathcal{X}^*} \|u - v\|_{\mathcal{X}}.$$

Combining this bound with (19) proves (9).

Now suppose $\omega_\theta(s) \ge m_\theta s^p$ for all $s \ge 0$, where $m_\theta > 0$ and $p > 1$. If $u = v$, the desired estimate is immediate. If $u \ne v$, then

$$m_\theta \|u - v\|_{\mathcal{X}}^p \le \left\| r^\theta \right\|_{\mathcal{X}^*} \|u - v\|_{\mathcal{X}}.$$

Dividing by the positive quantity $\|u - v\|_{\mathcal{X}}$ gives

$$m_\theta \|u - v\|_{\mathcal{X}}^{p-1} \le \left\| r^\theta \right\|_{\mathcal{X}^*},$$

and taking the $(p-1)$-st root completes the proof. □

Remark 2. *Unlike the contraction route, the preceding argument does not require an explicit fixed-point iteration. Once existence is supplied by the deterministic model theorem, the error estimate uses only the continuous residual and its stability modulus.*

Lemma 2 (Conditioning stability)**.** *Let $P, Q \in \mathcal{P}(\Theta)$ and $\Gamma \in \mathcal{B}(\Theta)$. If $\rho := \mathrm{TV}(P, Q)$ and $P(\Gamma) \ge \alpha > \rho$, then $Q(\Gamma) \ge \alpha - \rho > 0$ and*

$$\mathrm{TV}(Q(\cdot \mid \Gamma), P(\cdot \mid \Gamma)) \le \frac{2\rho}{\alpha - \rho}.$$

F.5 Proof of Lemma 2

Proof of Lemma 2. Let $\rho = \mathrm{TV}(P, Q)$ and assume $P(\Gamma) \ge \alpha > \rho$. By the definition of total variation, applied to the measurable set $A = \Gamma$,

$$|P(\Gamma) - Q(\Gamma)| \le \sup_{A \in \mathcal{B}(\Theta)} |P(A) - Q(A)| = \rho.$$

Therefore

$$Q(\Gamma) \ge P(\Gamma) - |P(\Gamma) - Q(\Gamma)| \ge P(\Gamma) - \rho \ge \alpha - \rho > 0.$$

Now fix an arbitrary measurable set $A \in \mathcal{B}(\Theta)$. Since both conditional probabilities are well defined,

$$\begin{aligned}&\left| \frac{Q(A \cap \Gamma)}{Q(\Gamma)} - \frac{P(A \cap \Gamma)}{P(\Gamma)} \right| \\ &\le \frac{|Q(A \cap \Gamma) - P(A \cap \Gamma)|}{Q(\Gamma)} + P(A \cap \Gamma) \left| \frac{1}{Q(\Gamma)} - \frac{1}{P(\Gamma)} \right| \\ &= \frac{|Q(A \cap \Gamma) - P(A \cap \Gamma)|}{Q(\Gamma)} + \frac{P(A \cap \Gamma)\,|P(\Gamma) - Q(\Gamma)|}{P(\Gamma) Q(\Gamma)}.\end{aligned}$$

Both absolute differences in the numerators are at most ρ by the definition of total variation, and $P(A \cap \Gamma) \le P(\Gamma)$. Hence

$$|Q(A \mid \Gamma) - P(A \mid \Gamma)| \le \frac{2\rho}{Q(\Gamma)} \le \frac{2\rho}{\alpha - \rho}.$$

Taking the supremum over all measurable sets $A \in \mathcal{B}(\Theta)$ proves the lemma. □

F.6 Proof of Theorem 5

Proof of Theorem 5. By the divergence-to-TV calibration (11),

$$\mathrm{TV}(Q_{N,y}, \Pi_y) \le \psi_D(D(Q_{N,y}\|\Pi_y)) \le \psi_D(\varepsilon_N) = r_N.$$

Because $r_N < \alpha$ and $\Pi_y(\Gamma) \ge \alpha$, Lemma 2 applies with $P = \Pi_y$ and $Q = Q_{N,y}$. Therefore

$$Q_{N,y}(\Gamma) \ge \alpha - r_N > 0$$

and

$$\mathrm{TV}(Q^{\Gamma}_{N,y}, \Pi^{\Gamma}_y) \le \frac{2r_N}{\alpha - r_N}.$$

The proposal-count statement follows from the geometric rejection-sampling calculation in Appendix F. □

F.7 Interpretation of the Upstream Posterior-Error Assumption

The bound $D(Q_{N,y}\|\Pi_y) \le \varepsilon_N$ in Theorem 5 is an upstream approximation assumption, not a quantity that certified SBI itself estimates from posterior samples. Its role is modular: whenever an SBI method, benchmark setting, or separate approximation analysis supplies such a bound (or any discrepancy that calibrates total variation), the theorem propagates it through certification. In applications where no quantitative upstream guarantee is available, retained mass and certificate margins remain observable diagnostics, but Theorems 5 and 6 should not be interpreted as data-driven estimates of an unknown divergence to the exact posterior.

F.8 Proof of Theorem 6

Proof of Theorem 6. Write $\beta = \Pi_y(\Gamma)$ and

$$\Delta_N = (1 - \beta) + \frac{2r_N}{\beta - r_N}.$$

By the triangle inequality,

$$\begin{aligned}\mathrm{TV}(Q^{\Gamma}_{N,y}, \Pi_y) &\le \mathrm{TV}(Q^{\Gamma}_{N,y}, \Pi^{\Gamma}_y) + \mathrm{TV}(\Pi^{\Gamma}_y, \Pi_y) \\ &\le \frac{2r_N}{\beta - r_N} + \mathrm{TV}(\Pi^{\Gamma}_y, \Pi_y),\end{aligned}$$

where Theorem 5 is applied with $\alpha = \beta$.

We compute the second term explicitly. For a measurable set A, let $x = \Pi_y(A \cap \Gamma)$ and $z = \Pi_y(A \cap \Gamma^c)$. Then $0 \le x \le \beta$, $0 \le z \le 1 - \beta$, and

$$\Pi^{\Gamma}_y(A) - \Pi_y(A) = \frac{x}{\beta} - x - z = \frac{1 - \beta}{\beta} x - z.$$

The two nonnegative terms on the right are each at most $1 - \beta$, so $\left|\Pi^{\Gamma}_y(A) - \Pi_y(A)\right| \le 1 - \beta$. Equality is attained for $A = \Gamma$ (equivalently $A = \Gamma^c$), hence

$$\mathrm{TV}(\Pi^{\Gamma}_y, \Pi_y) = 1 - \beta.$$

Therefore

$$\mathrm{TV}(Q^{\Gamma}_{N,y}, \Pi_y) \le \Delta_N.$$

Since total variation is at most one, the sharper displayed numerical bound is $\min\{1, \Delta_N\}$; the theorem is quantitatively informative when $\Delta_N < 1$.

For bounded measurable h, the variational identity associated with our TV convention is

$$\sup_{\|g\|_\infty \le 1} \left| \int_\Theta g \, d(P - Q) \right| = 2\,\mathrm{TV}(P, Q).$$

Taking $g = h / \|h\|_\infty$ when $h \not\equiv 0$ gives

$$\left| \mathbb{E}_{Q^{\Gamma}_{N,y}}[h] - \mathbb{E}_{\Pi_y}[h] \right| \le 2 \|h\|_\infty \mathrm{TV}(Q^{\Gamma}_{N,y}, \Pi_y) \le 2 \|h\|_\infty \Delta_N.$$

The case $h \equiv 0$ is immediate. This proves the observable bound. □

F.9 Downstream Decision-Risk Consequence

Derivation of the decision-risk bound in Theorem 6. For every action $a \in \mathcal{A}$, set $h_a(\theta) = \ell(a, \theta)$. Since $0 \le \ell(a, \theta) \le L_{\max}$, the preceding observable bound yields

$$\left| R_{Q^{\Gamma}_{N,y}}(a) - R_{\Pi_y}(a) \right| \le 2L_{\max}\Delta_N.$$

Assume first that minimizers exist and choose

$$\widehat{a} \in \arg\min_{a \in \mathcal{A}} R_{Q^{\Gamma}_{N,y}}(a), \qquad a^\star \in \arg\min_{a \in \mathcal{A}} R_{\Pi_y}(a).$$

Insert and subtract the certified-posterior risks:

$$\begin{aligned} R_{\Pi_y}(\widehat{a}) - R_{\Pi_y}(a^\star) = &\left[R_{\Pi_y}(\widehat{a}) - R_{Q^{\Gamma}_{N,y}}(\widehat{a})\right] \\ &+ \left[R_{Q^{\Gamma}_{N,y}}(\widehat{a}) - R_{Q^{\Gamma}_{N,y}}(a^\star)\right] \\ &+ \left[R_{Q^{\Gamma}_{N,y}}(a^\star) - R_{\Pi_y}(a^\star)\right]. \end{aligned}$$

The middle bracket is nonpositive by optimality of $\widehat{a}$. Each outer bracket is at most $2L_{\max}\Delta_N$, so

$$R_{\Pi_y}(\widehat{a}) - \inf_{a \in \mathcal{A}} R_{\Pi_y}(a) \le 4L_{\max}\Delta_N.$$

If a minimizer does not exist, take ε-minimizers for the two risks; the same calculation gives an additional 2ε, and letting $\varepsilon \downarrow 0$ yields the stated infimum form. □

F.10 Rejection-Cost Derivation

Derivation of the proposal-count bound. The proof of Theorem 5 already shows

$$a := Q_{N,y}(\Gamma) \ge \alpha - r_N > 0.$$

Under independent rejection sampling, each proposal is accepted with probability a. Hence the number K of proposals required for the first acceptance has the geometric distribution

$$\mathbb{P}(K = k) = (1 - a)^{k-1}a, \qquad k = 1, 2, \ldots.$$

Its expectation is

$$\mathbb{E}[K] = \sum_{k=1}^{\infty} k(1 - a)^{k-1}a = \frac{1}{a} \le \frac{1}{\alpha - r_N}.$$

This proves the claim. □

F.11 Runtime and Proposal-Cost Derivation

Let

$$a_{N,y} := Q_{N,y}(\Gamma)$$

be the probability that an independent posterior proposal passes the certificate. If generating one proposal costs C_{draw} and evaluating its certificate costs C_{cert}, the number of proposals required for one accepted draw is geometric with mean $1/a_{N,y}$. Hence

$$\mathbb{E}[\text{cost per certified draw}] = \frac{C_{\mathrm{draw}} + C_{\mathrm{cert}}}{a_{N,y}} \le \frac{C_{\mathrm{draw}} + C_{\mathrm{cert}}}{\alpha - r_N}, \tag{20}$$

where the last inequality follows from Theorem 5. Thus the computational cost has two distinct components: the cost of verifying each proposal and the number of proposals that must be checked before one is accepted.

Route I benchmark. For the controlled semilinear experiment, once $\theta = (A, \lambda)$ has been drawn, the parameter-level certificate only evaluates

$$q_\theta = 4\lambda T A(1 + 0.35)$$

and checks whether $q_\theta < 1$. This requires a fixed number of scalar arithmetic operations and no additional PDE solve, so $C_{\mathrm{cert}} = O(1)$ for this certificate. At the reported retained mass $a_{N,y} = 0.711$, one certified draw requires

$$\frac{1}{a_{N,y}} = 1.41$$

posterior proposals on average. In this benchmark, therefore, the additional sampling cost is determined primarily by the retained posterior mass rather than by certificate evaluation.

Route II finite-element verification. The cost can be substantially different when certification requires a continuous dual residual. Consider a Route II approximation with n_h degrees of freedom in a Hilbert-space setting. Let r_h be the assembled discrete residual vector and K_h the stiffness matrix associated with the Riesz map. The discrete dual residual norm is obtained by solving

$$K_h z_h = r_h, \qquad \eta_h := \|r_h\|_{H^{-1},h} = (r_h^\top K_h^{-1} r_h)^{1/2}.$$

A theorem-backed certificate additionally requires a reliable lifting bound from the discrete residual to the continuous residual, for example

$$\left\|r^\theta\right\|_{H^{-1}} \le \eta_h + \varepsilon_{\mathrm{lift}}(h,\theta).$$

Accordingly, we decompose the per-proposal verification cost as

$$C_{\mathrm{cert}}(n_h) = C_{\mathrm{assemble}}(n_h) + C_{\mathrm{Riesz}}(n_h) + C_{\mathrm{lift}}(n_h). \tag{21}$$

Here C_{assemble} is the cost of forming the residual vector (and the stiffness matrix when it is not already available), C_{Riesz} is the cost of solving $K_h z_h = r_h$, and C_{lift} is the cost of computing the reliable estimator $\varepsilon_{\mathrm{lift}}$ needed to control residual components not captured by the discrete space. The last term is essential because the Route II theorem requires a continuous dual norm rather than only a discrete algebraic residual.

The precise complexity depends on the discretization and linear solver. For a fixed-order finite-element method on a shape-regular mesh, residual assembly and standard local residual estimators require a number of local operations proportional to the number of degrees of freedom, so that

$$C_{\mathrm{assemble}}(n_h) = O(n_h), \qquad C_{\mathrm{lift}}(n_h) = O(n_h)$$

when a local reliable estimator is available. If the Riesz problem is solved by an iterative sparse method requiring N_{it} iterations, with $O(n_h)$ work per iteration, then

$$C_{\mathrm{Riesz}}(n_h) = O(n_h N_{\mathrm{it}}),$$

and therefore

$$C_{\mathrm{cert}}(n_h) = O(n_h(1+N_{\mathrm{it}}))\,. \tag{22}$$

With an optimal preconditioner or multigrid scheme for which $N_{\mathrm{it}} = O(1)$ under mesh refinement, this reduces to $C_{\mathrm{cert}}(n_h) = O(n_h)$. These rates are conditional rather than universal: a direct linear solver or a global verified lifting procedure can have a different complexity.

Combining (20) and (22), the corresponding expected cost per certified draw is, under the preceding finite-element assumptions,

$$\mathbb{E}[\text{cost per certified draw}] = O\left(\frac{C_{\mathrm{draw}}(n_h) + n_h(1+N_{\mathrm{it}})}{a_{N,y}}\right).$$

This expression separates the two computational bottlenecks of certified SBI: a small retained mass $a_{N,y}$ increases the number of proposals that must be verified, whereas an expensive continuous-residual calculation increases the cost of verifying each proposal.

G From Solver Output to a Theorem-Ready Certificate

The residual-to-error theorems in the main text are statements about continuous-model defects, whereas numerical solvers return finite-dimensional objects such as grid values, coefficient vectors, collocation losses, or discrete residuals. To apply the theorems, these numerical quantities must therefore be converted into rigorous upper bounds for the corresponding continuous defect. This section makes that conversion explicit in the settings used in this paper and distinguishes derived bounds from solver-dependent error estimates that must be supplied separately.

G.1 Route I: a computable defect bound for the controlled benchmark

For the semilinear benchmark, let

$$c := A(1+0.35), \qquad q_\theta := 4\lambda T c.$$

Appendix I proves that $q_\theta < 1$ implies both the self-map and contraction conditions, with contraction factor $\lambda_\theta = q_\theta$.

Let

$$v_K(t,x) = \sum_{|k|\le K} \widehat{v}_k(t) e^{2\pi i k x}$$

be a Fourier approximation. Its continuous Duhamel defect is

$$\mathcal{E}_\theta(v_K)(t) = v_K(t) - e^{0.08t\partial_x^2} u_0^A - \lambda \int_0^t e^{0.08(t-s)\partial_x^2} v_K(s)^2 \, ds.$$

Write the Fourier expansion of this defect as

$$\mathcal{E}_\theta(v_K)(t,x) = \sum_{k\in\mathbb{Z}} \widehat{\mathcal{E}}_{\theta,k}(t) e^{2\pi i k x},$$

which defines the coefficient $\widehat{\mathcal{E}}_{\theta,k}(t)$.

Suppose that the numerical implementation evaluates the defect at uniformly spaced times

$$t_n = n\Delta t, \qquad n = 0, \dots, N_t, \qquad \Delta t := T/N_t.$$

Thus Δt is the temporal sampling step used when evaluating the defect. For the retained Fourier modes, define

$$\delta_{\mathrm{grid}} := \max_{0\le n\le N_t} \sum_{|k|\le K} \left| \widehat{\mathcal{E}}_{\theta,k}(t_n) \right|.$$

Because

$$\left\| \sum_{|k|\le K} a_k e^{2\pi i k x} \right\|_{L_x^\infty} \le \sum_{|k|\le K} |a_k|,$$

δ_{grid} is an upper bound for the retained-mode defect at the sampled times.

To pass from this finite calculation to the continuous norm, three additional quantities must be controlled. Let $\varepsilon_{\mathrm{tail}}$ be a verified bound for the contribution of omitted Fourier modes, and let $\varepsilon_{\mathrm{quad}}$ be a verified bound for the error incurred when evaluating the Duhamel integral numerically. Finally, provided that the defect is differentiable in time with values in $L^\infty(\Omega)$, define

$$M_t := \sup_{0\le t\le T} \left\| \partial_t \mathcal{E}_\theta(v_K)(t) \right\|_{L^\infty}.$$

The argument below is useful only when this quantity is finite and a computable upper bound for it is available. We do not infer such a bound merely from time-grid values.

For any $t \in [0,T]$, choose a grid point t_n with $|t - t_n| \le \Delta t/2$. The fundamental theorem of calculus gives

$$\left\| \mathcal{E}_\theta(v_K)(t) - \mathcal{E}_\theta(v_K)(t_n) \right\|_{L^\infty} \le \int_{t_n}^t \left\| \partial_s \mathcal{E}_\theta(v_K)(s) \right\|_{L^\infty} ds \le \frac{\Delta t}{2} M_t.$$

Combining this estimate with the Fourier triangle inequality and the verified tail and quadrature bounds yields

$$\left\| \mathcal{E}_\theta(v_K) \right\|_{\mathcal{X}} \le \delta_{\mathrm{grid}} + \varepsilon_{\mathrm{tail}} + \varepsilon_{\mathrm{quad}} + \frac{\Delta t}{2} M_t =: \overline{\delta}_\theta. \tag{23}$$

Thus (23) is not an assumption once its four right-hand terms have been rigorously bounded; it follows from the triangle inequality and time interpolation above. Theorem 3 then gives

$$\left\| u^\theta - v_K \right\|_{\mathcal{X}} \le \frac{\overline{\delta}_\theta}{1 - q_\theta}. \tag{24}$$

The numerical verification supplies the numerator $\overline{\delta}_\theta$, while the deterministic contraction theorem supplies the amplification factor $(1 - q_\theta)^{-1}$.

G.2 Route II: from a discrete residual to a continuous dual residual

Consider first the Hilbert-space case

$$\mathcal{X} = H_0^1(\Omega), \qquad \|\phi\|_{\mathcal{X}} = \|\nabla\phi\|_{L^2},$$

and let

$$r^\theta := \mathcal{R}_\theta(v^\theta) \in H^{-1}(\Omega).$$

Its Riesz representative $z^\theta \in H_0^1(\Omega)$ is defined by

$$\int_\Omega \nabla z^\theta \cdot \nabla\phi \, dx = \langle r^\theta, \phi\rangle_{H^{-1},H_0^1} \qquad \text{for every } \phi \in H_0^1(\Omega).$$

Consequently,

$$\left\|r^\theta\right\|_{H^{-1}} = \left\|\nabla z^\theta\right\|_{L^2}.$$

Let $\mathcal{X}_h \subset H_0^1(\Omega)$ be a finite-element space, let r_h be the assembled residual vector, and let K_h be the stiffness matrix. The discrete Riesz representative solves

$$K_h z_h = r_h,$$

and the corresponding discrete dual norm is

$$\eta_h := (r_h^\top K_h^{-1} r_h)^{1/2}.$$

Importantly, η_h alone controls only the residual tested against $\mathcal{X}_h$. It is not automatically equal to the continuous H^{-1} norm.

A theorem-backed Route II certificate therefore requires a *reliable residual estimator*: a computable quantity $\varepsilon_{\mathrm{lift}}(h,\theta) \geq 0$ satisfying

$$\left\|r^\theta\right\|_{H^{-1}} \leq \eta_h + \varepsilon_{\mathrm{lift}}(h,\theta). \tag{25}$$

The form of $\varepsilon_{\mathrm{lift}}$ is discretization dependent. For a standard conforming finite-element method, it may be obtained from element residuals and interelement flux jumps through a reliable a posteriori estimator (Repin, 2008). We deliberately do not assign a universal formula to $\varepsilon_{\mathrm{lift}}$ because such a formula would be false across different PDEs and discretizations.

Once a reliable bound of the form (25) has been established, strong monotonicity with constant $m_\theta > 0$ gives

$$\left\|u^\theta - v_h^\theta\right\|_{\mathcal{X}} \leq \frac{\eta_h + \varepsilon_{\mathrm{lift}}(h,\theta)}{m_\theta}. \tag{26}$$

In the controlled Route II experiment of Appendix I.3, no lifting estimator is needed because the continuous H^{-1} residual is evaluated analytically. There $m_\theta = 1$, and the resulting effectivity is 1.23–1.28.

G.3 How numerical error budgets enter the certificate

The frequently used decomposition

$$d_{\mathrm{cont}} \leq d_h + \varepsilon_{\mathrm{disc}} + \varepsilon_{\mathrm{quad}} + \varepsilon_{\mathrm{round}}$$

is not a universal numerical-analysis theorem by itself. It is shorthand for a triangle-inequality argument that is valid only after the corresponding approximations have been defined.

To make this precise, let $D_\theta(v)$ denote the exact continuous defect and let $\widehat{D}_{\theta,h}(v)$ denote the quantity actually evaluated by the numerical implementation. Introduce intermediate approximations so that

$$D_\theta(v) - \widehat{D}_{\theta,h}(v) = E_{\mathrm{disc}} + E_{\mathrm{quad}} + E_{\mathrm{round}},$$

where E_{disc} denotes the error caused by finite-dimensional discretization or truncation, E_{quad} the error in numerical integration or time stepping, and E_{round} the floating-point evaluation error. If verified estimates

$$\|E_{\mathrm{disc}}\| \leq \varepsilon_{\mathrm{disc}}, \qquad \|E_{\mathrm{quad}}\| \leq \varepsilon_{\mathrm{quad}}, \qquad \|E_{\mathrm{round}}\| \leq \varepsilon_{\mathrm{round}}$$

are available, then the triangle inequality gives

$$\|D_\theta(v)\| \le \left\|\widehat{D}_{\theta,h}(v)\right\| + \varepsilon_{\mathrm{disc}} + \varepsilon_{\mathrm{quad}} + \varepsilon_{\mathrm{round}}. \tag{27}$$

Thus, in the notation sometimes used for brevity, $d_{\mathrm{cont}} := \|D_\theta(v)\|$ and $d_h := \left\|\widehat{D}_{\theta,h}(v)\right\|$. Equation (23) is a concrete spectral instance of this general triangle-inequality construction.

The individual error terms are solver dependent. A finite-element method may use residual and flux-jump estimators for $\varepsilon_{\mathrm{disc}}$; a spectral method may use a Fourier-tail estimate; and verified quadrature or interval arithmetic can be used when integration or rounding errors are non-negligible. The essential requirement is that each term be an upper bound in the same continuous norm required by the deterministic theorem.

G.4 PINNs: from collocation loss to a continuous residual

Let D be a compact space–time domain and let

$$z_i \in D, \qquad i = 1, \dots, N_{\mathrm{coll}},$$

be the collocation points used to evaluate the PINN residual. Here N_{coll} is the total number of collocation points. Define their fill distance by

$$h_{\mathrm{fill}} := \sup_{z \in D} \min_{1 \le i \le N_{\mathrm{coll}}} \|z - z_i\| .$$

Thus h_{fill} measures the largest distance from a point of the continuous domain to the nearest collocation point.

Let $\rho_\theta : D \to \mathbb{R}$ denote the differential residual of the learned state. Assume that ρ_θ is Lipschitz on D with a known upper bound L_ρ on its Lipschitz constant:

$$|\rho_\theta(z) - \rho_\theta(z')| \le L_\rho \|z - z'\| \qquad (z, z' \in D).$$

For any $z \in D$, choose a nearest collocation point z_i. Then

$$|\rho_\theta(z)| \le |\rho_\theta(z_i)| + L_\rho h_{\mathrm{fill}},$$

and therefore

$$\|\rho_\theta\|_{L^\infty(D)} \le \max_{1 \le i \le N_{\mathrm{coll}}} |\rho_\theta(z_i)| + L_\rho h_{\mathrm{fill}}.$$

If the reported collocation loss is

$$L_{\mathrm{coll}} = \frac{1}{N_{\mathrm{coll}}} \sum_{i=1}^{N_{\mathrm{coll}}} |\rho_\theta(z_i)|^2,$$

then

$$\max_i |\rho_\theta(z_i)| \le \left(\sum_{i=1}^{N_{\mathrm{coll}}} |\rho_\theta(z_i)|^2 \right)^{1/2} = \sqrt{N_{\mathrm{coll}} L_{\mathrm{coll}}},$$

which yields

$$\|\rho_\theta\|_{L^\infty(D)} \le \sqrt{N_{\mathrm{coll}} L_{\mathrm{coll}}} + L_\rho h_{\mathrm{fill}}. \tag{28}$$

This bound also makes its limitation explicit. A small mean collocation loss does not by itself imply a small continuous residual: one must additionally control both the spacing of the collocation points and the regularity of the residual. The stress test in Appendix I.4 demonstrates what can fail when this lifting information is absent.

For an evolution equation, one further converts the differential residual into the Route I fixed-point defect. If

$$\rho_\theta = \partial_t v - A_\theta v - F_\theta(v), \qquad e_0 = v(0) - u_0,$$

and

$$\|S_\theta(t)\| \le M_\theta e^{\omega_\theta t}, \qquad 0 \le t \le T,$$

then variation of constants gives, for the corresponding supremum-type state norm,

$$\|\mathcal{E}_\theta(v)\|_{\mathcal{X}} \le M_\theta e^{\omega_\theta T} \left(\|e_0\| + T \|\rho_\theta\|_{L^\infty(D)} \right).$$

Combining this estimate with (28) yields a computable Route I defect bound whenever L_ρ, h_{fill}, and the linear-evolution constants are themselves rigorously bounded.

H Scientific Model Instantiations: What Changes and What Does Not

The purpose of this section is not to catalogue PDE function spaces. It shows how the same certified-SBI layer is instantiated in model classes that are common in scientific inference. Across all examples, the learned posterior and conditioning operation are unchanged; what changes is the deterministic theorem that defines $C(\theta)$ and the norm in which the state is verified.

I Experimental Protocols and Additional Results

This section supplies the numerical details behind the main experiment and additional checks of the two certification routes. We distinguish a controlled Route I posterior benchmark, a controlled Route II residual-to-error experiment, a collocation stress test, and an incompressible-flow deployment protocol. Quantitative claims are made only for the controlled experiments.

I.1 Model-Specific Quantities and Costs

Table 2: Continuous-model quantities required by representative certificates.

Model	Theorem-ready quantity	Main verification cost
Semilinear heat	$\|\mathcal{E}_\theta(v)\|_{\mathcal{X}}/(1-\lambda_\theta)$	space–time defect and quadrature
Incompressible flow	Duhamel defect in the selected flow norm	nonlinear convolution and continuous-norm bound
Nonlinear steady state	$\|\mathcal{A}_\theta(v)-f_\theta\|_{\mathcal{X}^*}$ plus stability modulus	dual residual / Riesz solve
Nonlinear conductivity	weak residual in $\mathcal{X}^*$ plus coercivity/monotonicity constant	verified quadrature and stability constant

These are continuous-model quantities rather than solver-native scores. Appendix G explains how numerical outputs are lifted to the required norms.

I.2 What Is Actually Verified Under Solver Refinement

For a retained parameter, let v_h^θ be an approximate state. A theorem-backed Route I bound first requires a computable continuous-defect upper bound

$$\|\mathcal{E}_\theta(v_h^\theta)\|_{\mathcal{X}} \le \overline{\delta}_h(\theta).$$

For the spectral benchmark, Appendix G gives

$$\overline{\delta}_h = \delta_{\mathrm{grid}} + \varepsilon_{\mathrm{tail}} + \varepsilon_{\mathrm{quad}} + \frac{\Delta t}{2} M_t.$$

Only after these four terms have been bounded do we define

$$b_h(\theta) := \frac{\overline{\delta}_h(\theta)}{1-\lambda_\theta}.$$

If an independently controlled reference solution u_{ref}^θ is available, $e_h(\theta) = \|v_h^\theta - u_{\mathrm{ref}}^\theta\|_{\mathcal{X}}$ can be compared with $b_h(\theta)$. The ratio b_h/e_h measures sharpness; it does not replace the proof that $\overline{\delta}_h$ bounds the continuous defect. In particular, visual mesh convergence is not used as a rigorous value of the truncation or quadrature terms.

I.3 Route II Verification Experiment: A Steady Nonlinear Model

This controlled experiment tests the part of certified SBI not exercised by the Route I posterior benchmark: after a Route II draw is accepted, does the continuous dual residual give a quantitatively useful error bar for an approximate state? This question is relevant to finite-element and neural variational solvers, which naturally return approximate equilibria.

Consider

$$-u'' + u^3 = f, \qquad u(0) = u(1) = 0, \tag{29}$$

on $\mathcal{X} = H_0^1(0,1)$ with $\|u\|_{\mathcal{X}} = \|u'\|_{L^2}$. Choose

$$u_\star(x) = \sin(\pi x), \qquad f(x) = \pi^2 \sin(\pi x) + \sin^3(\pi x),$$

so that $u_\star$ is the exact solution. For $\mathcal{A}(u) = -u'' + u^3$,

$$\begin{aligned}\langle \mathcal{A}(u) - \mathcal{A}(v), u - v \rangle_{H^{-1}, H_0^1} &= \|u - v\|_{\mathcal{X}}^2 \\ &\quad + \int_0^1 (u^3 - v^3)(u - v)\, dx \\ &\geq \|u - v\|_{\mathcal{X}}^2,\end{aligned}$$

because $(u^3 - v^3)(u - v) = (u - v)^2(u^2 + uv + v^2) \geq 0$. Thus the stability modulus is $\omega(s) = s^2$, and Theorem 4 gives

$$\|u_\star - v\|_{\mathcal{X}} \leq \|\mathcal{A}(v) - f\|_{H^{-1}}.$$

For $v_\delta = (1 + \delta)u_\star$, set $b_\delta = (1 + \delta)^3 - 1$. Then

$$r_\delta = \mathcal{A}(v_\delta) - f = \left(\delta\pi^2 + \frac{3b_\delta}{4}\right) \sin(\pi x) - \frac{b_\delta}{4} \sin(3\pi x).$$

Orthogonality of the sine basis gives

$$e_\delta := \|u_\star - v_\delta\|_{\mathcal{X}} = \frac{|\delta|\pi}{\sqrt{2}},$$

and

$$B_\delta := \|r_\delta\|_{H^{-1}}, \qquad B_\delta^2 = \frac{\left(\delta\pi^2 + 3b_\delta/4\right)^2}{2\pi^2} + \frac{b_\delta^2}{288\pi^2}.$$

Thus e_δ is the exact state error and B_δ is the exact continuous Route II upper bound; no reference discretization is required.

To quantify the sharpness of the certificate, define the *effectivity index*

$$\mathrm{Eff}_\delta := \frac{B_\delta}{e_\delta}.$$

The residual-to-error theorem implies $\mathrm{Eff}_\delta \geq 1$ for $\delta \neq 0$; values close to one indicate that the certified upper bound is quantitatively sharp.

Table 3: **Route II residual-to-error verification.** Here $e_\delta = \|u_\star - v_\delta\|_{\mathcal{X}}$ is the exact state error, $B_\delta = \|\mathcal{A}(v_\delta) - f\|_{H^{-1}}$ is the exact continuous Route II upper bound, and $\mathrm{Eff}_\delta = B_\delta / e_\delta$ measures its sharpness.

δ	e_δ	B_δ	Eff_δ
0.02	0.0444	0.0548	1.23
0.05	0.1111	0.1377	1.24
0.10	0.2221	0.2781	1.25
0.20	0.4443	0.5673	1.28

The effectivity index remains between 1.23 and 1.28 over the tested perturbation range. Thus, in this controlled example, the Route II residual bound overestimates the exact state error by only about 23–28%. This supports the state-level claim underlying Route II: quantitative monotonicity and a continuous dual residual can provide an informative error bar for an approximate steady state. It does not test posterior conditioning; the Route I benchmark serves that separate purpose.

I.4 Why Finite Collocation Loss Is Not a Continuous Certificate

This stress test explains why the paper distinguishes a solver training diagnostic from the continuous residual required by Theorem 4. For (29), let

$$v_k(x) = u_\star(x) + \frac{\delta}{k} \sin(k\pi x)$$

and use collocation points $x_j = j/k$, $j = 0, \dots, k$. At these points, $v_k(x_j) = u_\star(x_j)$ and the oscillatory contribution to $-v_k''$ also vanishes, so the pointwise PDE residual is zero at every collocation point. Nevertheless,

$$\|v_k - u_\star\|_{\mathcal{X}} = \frac{|\delta|\pi}{\sqrt{2}},$$

independent of k. For $\delta = 0.05$ this error is 0.1111.

There is no contradiction: finitely many point evaluations do not determine a continuous residual norm, and the oscillatory mode lies between sampled locations. Thus a PINN or other collocation solver can have small empirical residual without satisfying the hypothesis of the certified residual-to-error theorem. Appendix G gives an explicit Lipschitz/fill-distance lifting bound that closes this gap when its assumptions can be verified. This example motivates that additional verification step; it is not a claim that collocation or PINN methods are intrinsically unreliable.

I.5 Controlled Semilinear Benchmark: Full Protocol

We solve (16) on the periodic unit interval. The ground truth is $(A, \lambda) = (0.75, 1.20)$, $T = 0.15$, and observations are taken at

$$(t, x) = (0.05, 0.10),\ (0.10, 0.35),\ (0.15, 0.70)$$

with independent Gaussian noise of standard deviation 0.025. The reference trajectory uses 96 Fourier modes and time step 2.5×10^{-4}; inference uses 32 modes and time step 10^{-3}. Random seed 27 is used. The prior is uniform on the rectangle stated in Section 5.

This is the Route I posterior experiment. Route II is tested separately in Appendix I.3; separating them makes clear which theoretical guarantee each experiment evaluates.

Write

$$T_\theta(u)(t) = e^{0.08t\partial_x^2} u_0 + \lambda \int_0^t e^{0.08(t-s)\partial_x^2} u(s)^2 \, ds.$$

Let $c = \|u_0\|_{L^\infty} = A(1 + 0.35)$ and choose $R_\theta = 2c$. Heat-flow contractivity gives

$$\|T_\theta(u)\|_{\mathcal{X}} \le c + \lambda T R_\theta^2 = c + 4\lambda T c^2.$$

Thus the ball is invariant if $4\lambda T c \le 1$. For u, v in this ball,

$$\begin{aligned} \|T_\theta(u) - T_\theta(v)\|_{\mathcal{X}} &\le \lambda T \|u^2 - v^2\|_{\mathcal{X}} \\ &\le \lambda T (\|u\|_{\mathcal{X}} + \|v\|_{\mathcal{X}}) \|u - v\|_{\mathcal{X}} \\ &\le 4\lambda T c \, \|u - v\|_{\mathcal{X}}. \end{aligned}$$

Hence

$$\lambda_\theta = 4\lambda T A(1 + 0.35), \qquad C(\theta) = \begin{cases} 1, & \lambda_\theta < 1, \\ 0, & \lambda_\theta \ge 1. \end{cases}$$

This case distinction spells out the $\{0, 1\}$-valued certificate indicator used in the main text; the plotted boundary is the contraction threshold $\lambda_\theta = 1$, not a heuristic tuning rule.

We also justify the TV identity used in the reported results. Let p be a probability measure, $a = p(\Gamma) > 0$, and $p^\Gamma = p(\cdot \mid \Gamma)$. For measurable A, put $x = p(A \cap \Gamma)$ and $z = p(A \cap \Gamma^c)$. Then

$$p^\Gamma(A) - p(A) = \frac{x}{a} - x - z = \frac{1 - a}{a} x - z.$$

Because $0 \le (1 - a)x/a \le 1 - a$ and $0 \le z \le 1 - a$,

$$|p^\Gamma(A) - p(A)| \le 1 - a.$$

Taking $A = \Gamma$ attains equality. Therefore, for $\mathrm{TV}(P, Q) = \sup_A |P(A) - Q(A)|$,

$$\mathrm{TV}(p, p^\Gamma) = 1 - p(\Gamma). \tag{30}$$

At retained mass 0.711, this gives 0.289, exactly the TV change reported in the main text.

Table 4: Certification-threshold ablation. Smaller τ imposes a stricter acceptance rule. TV change follows from (30); proposals/accept is the reciprocal of retained mass.

τ	Retained mass	TV change	Proposals/accept
0.70	0.266	0.734	3.75
0.85	0.405	0.595	2.47
1.00	0.711	0.289	1.41

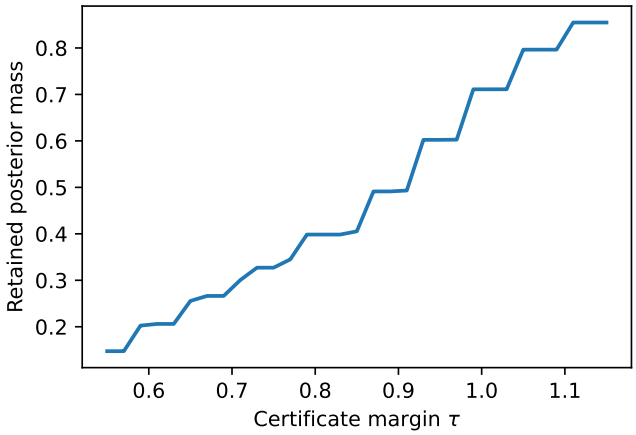


Figure 3: **Coverage–cost effect of the certification threshold.** The horizontal axis is the threshold τ in $\Gamma_\tau = \{\lambda_\theta < \tau\}$; the vertical axis is the learned posterior mass $Q_{N,y}(\Gamma_\tau)$ retained after certification. The curve must increase with τ because the accepted sets are nested. Moving left imposes a stricter certificate, reducing coverage and increasing the expected rejection cost $1/Q_{N,y}(\Gamma_\tau)$.

I.6 Certification-Threshold Ablation

For $\tau \in (0, 1]$, define the stricter diagnostic set

$$\Gamma_\tau = \{\theta : \lambda_\theta < \tau\}.$$

Smaller τ requires a draw to lie farther inside the contraction regime. If $\tau_1 < \tau_2$, then $\Gamma_{\tau_1} \subseteq \Gamma_{\tau_2}$, so $Q_{N,y}(\Gamma_\tau)$ is necessarily nondecreasing in τ. A monotonically increasing retained-mass curve is therefore the expected mathematical behavior.

This ablation does not by itself prove that one threshold is preferable. Instead, it supports the paper's diagnostic claim: certificate strictness has a measurable coverage and computational price, and retained mass makes that price visible.

I.7 Incompressible Navier–Stokes: Deployment Protocol

This subsection concerns the incompressible Navier–Stokes equation, not a compressible-flow model. On a periodic domain it has the form

$$\partial_t u - \nu\Delta u + (u \cdot \nabla)u + \nabla p = 0, \qquad \nabla \cdot u = 0.$$

The divergence-free constraint removes density as an independent dynamical variable and allows the pressure to be eliminated with the Leray projector. A compressible model contains additional density/thermodynamic variables and requires a different deterministic well-posedness theorem; it is outside the present certificate.

For a finite-dimensional test, parameterize a divergence-free initial velocity by Fourier coefficients,

$$u_0^\xi(x) = \sum_{k \in \mathcal{K}} \xi_k \frac{k^\perp}{|k|} e^{ik\cdot x}, \qquad \xi_{-k} = \overline{\xi_k}.$$

The linear heat evolution is explicit mode by mode, while the nonlinear term is evaluated through Fourier convolution. A Route I certificate would require a verified upper bound in the same continuous flow norm used by the selected Navier–Stokes theorem, including Fourier truncation, time quadrature, supremum approximation, and rounding errors.

We include this as a deployment protocol rather than a completed experiment: the current paper does not report a theorem-backed Navier–Stokes posterior result. Its role is to show how the abstract Route I interface would be instantiated for a canonical fluid model and to identify the verification budget that remains before such a result can be claimed.

J Extended Limitations and Future Directions

A certificate is sufficient, not necessary. For the certificates used here, $C(\theta) = 1$ proves that the assumptions of a specified deterministic theorem hold. The converse is not asserted: $C(\theta) = 0$ does not imply nonexistence, nonuniqueness, or physical invalidity. A parameter rejected by one theorem may be covered by a sharper theorem or by a different solution class. If several measurable certificate indicators C_j are available, their certified sets $\Gamma_j = \{\theta : C_j(\theta) = 1\}$ may be combined through $\Gamma = \bigcup_j \Gamma_j$; finite or countable unions remain measurable. Care is still required to record which theorem supplies the guarantee for each accepted draw.

Numerical verification can be the dominant cost. The abstract residual-to-error inequalities become computational certificates only after all quantities in them have rigorous upper or lower bounds. Appendix G shows explicitly where truncation, quadrature, time interpolation, residual lifting, and rounding errors enter. For the controlled Route I benchmark the certificate query is only a scalar inequality, but for a Route II finite-element state it can require residual assembly, a Riesz solve, and a reliable lifting estimator. Thus certificate cost is application dependent and can exceed the cost of an empirical solver diagnostic.

Conditioning deliberately changes the target. Certified SBI reports the conditional distribution $\Pi_y^\Gamma = \Pi_y(\cdot \mid \Gamma)$, not the unrestricted posterior Π_y. Equation (30) shows that the change caused by exact hard conditioning is exactly $1 - \Pi_y(\Gamma)$ in total variation. Consequently, if the target posterior places little mass on Γ, no post-processing argument can make the certified and unrestricted targets close. This is why retained mass is a primary diagnostic rather than an implementation detail.

Sharper and adaptive certificates are future work. The current experiments use fixed sufficient conditions. A sharper local certificate—for example, a verified Newton–Kantorovich, inf–sup, or computer-assisted argument centered at an individual numerical state—could enlarge the certified region without retraining the SBI estimator. A separate research direction is to use certificate information during sequential simulation allocation. Such a method would need to preserve the statistical corrections required by the chosen SBI algorithm; simply steering simulations toward easily certified parameters would in general change the training distribution and is not analyzed in this paper.

Extension to real scientific data. The quantitative experiments in this paper use controlled problems for which the continuous-model certificate and its error budget can be evaluated explicitly. Extending certified SBI to experimental measurements is an important next step rather than a result claimed here. For example, the open 2D electrical-impedance-tomography data of Hauptmann et al. (2017) contain laboratory current-injection and boundary-voltage measurements and provide a natural Route II conductivity setting. A theorem-backed deployment would additionally need to incorporate geometry and boundary/electrode modeling, measurement preprocessing, and discretization error into the same continuous residual or stability bound used by the certificate. Developing such end-to-end real-data certificates is left for future work.

K What Counts as Numerical Certification, and How to Reproduce It

This distinction is important because the main theorems are statements about continuous models, whereas experiments are performed with finite numerical representations. We use *theorem-backed numerical certificate* to mean that every numerical quantity entering the deterministic theorem has been replaced by a verified bound in the required continuous norm. Depending on the solver, this can require: (i) verified stability or contraction constants; (ii) a continuous-space residual/defect bound; (iii) truncation and interpolation estimates; (iv) quadrature or time-integration error; (v) verified control of any supremum used in the norm; and (vi) rounding-error control when it is not negligible (Repin, 2008; Nakao et al., 2019). A small empirical residual or a visually converged mesh is therefore a useful diagnostic, but it is not by itself the mathematical certificate asserted by the theorems.

Reproducibility. To facilitate reproduction of the reported experiments, we specify the parameter domain, ground-truth parameter, observation locations and noise level, random seed, Fourier resolutions, time steps, certificate inequality, and posterior-conditioning rule for the controlled semilinear benchmark in Appendix I. For the controlled Route II experiment, the exact solution and continuous residual are given analytically, allowing the reported state errors and effectivity indices to be computed directly.

Reproducibility and rigorous numerical certification are distinct. A numerical result may be reproducible without being theorem-backed if, for example, its truncation or quadrature error has not been rigorously bounded. Conversely, a theorem-backed numerical certificate must specify the numerical procedure and error bounds used to evaluate the quantities entering the certificate. The experiments in this paper are labeled accordingly: quantitative certification claims are made only where the required continuous-model bounds are established.