

Expert Systems With Applications

Gated Physics-Informed Neural Networks for Inverse Local Volatility Calibration with Learned Regime Structure

--Manuscript Draft--

Manuscript Number:	
Article Type:	Full length article
Section/Category:	1.1 AI / ML model and method
Keywords:	Physics-informed neural networks; Inverse problems; Local volatility calibration; Mixture of experts; Regime decomposition; Financial PDEs
Corresponding Author:	Manoov Rajapandy Vellore Institute of Technology Vellore, Tamil Nadu INDIA
First Author:	Manoov Rajapandy
Order of Authors:	Manoov Rajapandy Vrinda Bajaj Veydant Katyal
Abstract:	Local volatility calibration from discrete option prices is inherently ill-posed, and the classical Dupire formula fails catastrophically under realistic market conditions of noise and sparsity. Physics-Informed Neural Networks offer a principled alternative through PDE-constrained learning, but standard joint optimisation of price and volatility networks encounters a fourth- order gradient pathology that limits convergence quality. We propose a decoupled PINN architecture in which the call-price surface is learned under Black-Scholes constraints whilst volatility is recovered via independent regression. A stop-gradient operation eliminates the gradient pathology by breaking the computational dependency of PDE residuals on volatility parameters. This architectural choice yields a structural invariance: volatility error becomes independent of the PDE loss weight λ_{PDE}, removing what is typically the most consequential hyperparameter in such frameworks. Additionally, we represent volatility as a gated mixture of two expert networks, enabling regime structure to emerge naturally during training. The method is validated on synthetic two-regime surfaces and real equity options, demonstrating substantial improvements in reconstruction stability under data sparsity compared to Dupire- based methods. No-arbitrage constraints are satisfied consistently throughout the calibrated domain.

Vrinda Bajaj, Veydant Katyal, Manoov Rajapandy

School of Computer Science and Engineering, Vellore Institute of Technology, Vellore, India

To Editor,

Journal of Expert Systems with Applications,

Elsevier.

We are submitting our manuscript, "Gated Physics-Informed Neural Networks for Inverse Local Volatility Calibration with Learned Regime Structure" for consideration as an original research paper.

We address the unresolved computational difficulty of inverse physics-informed neural networks for local volatility calibration, where jointly training the price and PDE-residual networks induces a fourth-order gradient pathology in the volatility field and destabilizes training under sparse or noisy data. Here we develop a decoupled training architecture that severs this pathological coupling with a stop-gradient operation between the price and volatility networks, paired with a mixture-of-experts gating mechanism, GatedSigmaNet, which resolves the volatility surface into continuously blended functional regimes rather than a single global fit.

The method's reliability is established by a controlled noise-and-sparsity robustness benchmark against the standard Dupire finite-difference baseline, multi-seed repetition, an ablation sweep across five orders of magnitude of the PDE weighting coefficient confirming structural λ -invariance of the volatility RMSE, and a real-data feasibility study on SPY options, with GatedPINN retaining a volatility R^2 of 0.97 at 10% data density against the Dupire baseline's 0.342 and RMSE reductions of $4.8\times$ to $38.9\times$ across tested conditions. This advance is needed because local volatility calibration from sparse, noisy option-price data is a classically ill-posed inverse problem that existing coupled PINN formulations cannot calibrate reliably once data density falls below levels common in practice.

We believe Journal of Expert Systems with Applications is the right Journal because our contribution is a general decoupling and gating methodology for physics-informed inverse

problems, verified through a reusable benchmark protocol, rather than an application of an existing method or a pure numerical-analysis result.

This manuscript is original, has not been published previously, and is not under consideration for publication elsewhere. All authors have read and approved the submission. We declare no competing interests.

Sincerely,
Dr. Manoov Rajapandy,
and co-authors Vrinda Bajaj and Veydant Katyal.

Gated Physics-Informed Neural Networks for Inverse Local Volatility Calibration with Learned Regime Structure

Vrinda Bajaj^{a,1}, Veydant Katyal^{a,1}, Manooov Rajapandy^{a,*}

^a*School of Computer Science and Engineering, Vellore Institute of Technology, Vellore, Tamil Nadu, India*

Abstract

Local volatility calibration from discrete option prices is inherently ill-posed, and the classical Dupire formula fails catastrophically under realistic market conditions of noise and sparsity. Physics-Informed Neural Networks offer a principled alternative through PDE-constrained learning, but standard joint optimisation of price and volatility networks encounters a fourth-order gradient pathology that limits convergence quality. We propose a decoupled PINN architecture in which the call-price surface is learned under Black-Scholes constraints whilst volatility is recovered via independent regression. A stop-gradient operation eliminates the gradient pathology by breaking the computational dependency of PDE residuals on volatility parameters. This architectural choice yields a structural invariance: volatility error becomes independent of the PDE loss weight λ_{PDE} , removing what is typically the most consequential hyperparameter in such frameworks. Additionally, we represent volatility as a gated mixture of two expert networks, enabling regime structure to emerge naturally during training. The method is validated on synthetic two-regime surfaces and real equity options, demonstrating substantial improvements in reconstruction stability under data sparsity compared to Dupire-based methods. No-arbitrage constraints are satisfied consistently throughout the calibrated domain.

Keywords: Physics-informed neural networks, Inverse problems, Local volatility calibration, Mixture of experts, Regime decomposition, Dupire equation, Decoupled optimisation, Financial PDEs

Graphical Abstract. Overview of the GatedPINN framework: market option prices (left) are processed through a decoupled inverse PINN that simultaneously trains PriceNet on the Black–Scholes PDE constraint and GatedSigmaNet on volatility regression targets, yielding a calibrated local volatility surface with interpretable regime decomposition (right).

*Corresponding author.

Email address: mailmanooov@gmail.com (Manooov Rajapandy)

¹These authors contributed equally to this work.

²These authors contributed equally to this work.

Gated Physics-Informed Neural Networks for Inverse Local Volatility Calibration with Learned Regime Structure

Abstract

Local volatility calibration from discrete option prices is inherently ill-posed, and the classical Dupire formula fails catastrophically under realistic market conditions of noise and sparsity. Physics-Informed Neural Networks offer a principled alternative through PDE-constrained learning, but standard joint optimisation of price and volatility networks encounters a fourth-order gradient pathology that limits convergence quality. We propose a decoupled PINN architecture in which the call-price surface is learned under Black-Scholes constraints whilst volatility is recovered via independent regression. A stop-gradient operation eliminates the gradient pathology by breaking the computational dependency of PDE residuals on volatility parameters. This architectural choice yields a structural invariance: volatility error becomes independent of the PDE loss weight λ_{PDE} , removing what is typically the most consequential hyperparameter in such frameworks. Additionally, we represent volatility as a gated mixture of two expert networks, enabling regime structure to emerge naturally during training. The method is validated on synthetic two-regime surfaces and real equity options, demonstrating substantial improvements in reconstruction stability under data sparsity compared to Dupire-based methods. No-arbitrage constraints are satisfied consistently throughout the calibrated domain.

Keywords: Physics-informed neural networks, Inverse problems, Local volatility calibration, Mixture of experts, Regime decomposition, Dupire equation, Decoupled optimisation, Financial PDEs


Graphical Abstract. Overview of the GatedPINN framework: market option prices (left) are processed through a decoupled inverse PINN that simultaneously trains PriceNet on the Black–Scholes PDE constraint and GatedSigmaNet on volatility regression targets, yielding a calibrated local volatility surface with interpretable regime decomposition (right).

*Corresponding author.

Email address: mailmanoov@gmail.com (Manoov Rajapandy)

¹These authors contributed equally to this work.

²These authors contributed equally to this work.



1. Introduction

1.1. Motivation and Context

The local volatility model developed by Dupire [1] and Derman and Kani [2] has a special place in quantitative finance. As the only continuous diffusion model which can be constructed from a complete, arbitrage-free set of option prices, it offers a model for a rich set of market data. The local volatility model relies on the forward Black-Scholes equation in the strike-maturity plane to connect the local volatility surface $\sigma(t, S)$ to the observable call price surface $C(\tau, K)$ through the ratio of first- and second-order partial derivatives. Given this model, calibration is considered to solve an inverse problem in which the local volatility surface is recovered from a finite and noisy data sample of option prices.

In a functional-analytic perspective, the inverse problem is ill-posed. The mapping of σ to prices is compact in the natural function spaces, and thus, the inverse is unbounded; the recovered volatility surface can have large spikes given the compact mapping of prices can have small changes. While Dupire's formula is an analytic solution, it requires the inversion to be implemented using the numerical differentiation of the total implied variance surface, which is interpolated and has the second derivative of C with respect to K , which is also an interpolated quantity. The estimates of local volatility, given the conditions of the market, are very noisy and therefore are of no use. These noisy conditions include the presence of a boundary of near expiration options, illiquid and distant options, and an overall lack of orders in the market. Dupire-calibrated volume surfaces are documented and described as needing a significant amount of editing after the initial calibration and often fail to yield an acceptable surface, even after the implementation of careful regularization, when the data is sparse.

Volatility surfaces exhibit the signature of fractal structures and do not show monotonic smooth transitions. Periods of calm intersperse with volatility spikes. The calm periods are characterized by levels of low volatility and symmetric smiles and are interrupted by crisis segments that contain surfaces that are skewed in a negative direction. The frequent appearance of these structures within empirical studies of volatility present a challenge for calibration methods.

1.2. Physics-Informed Neural Networks for Computational Finance

Physics-Informed Neural Networks (PINNs) [14] strike a balance between data-centric regression techniques and the classical finite-difference method. Since the training loss of PINNs is a function of the residuals of PDEs, it defines a surface within a constraining manifold, where some form of implicit regularization is built-in. This characteristic is beneficial when the data is limited or of low quality. The utility of PINNs has been demonstrated in several studies across fluid and solid mechanics [15, 16] and in the biomedical domain [17]. In the context of computational finance, the parabolic equations of the Black-Scholes framework define the natural problem domain. In the context of the forward problem of pricing certain families of parametric options, the works of Salvador et al. [18] and Dhiman and Hu [19] achieved a level of accuracy comparable to available closed-form solutions. For the inverse problem, recovering $\sigma(t, S)$ from observed prices, a growing literature has

emerged: Berardi, Cosma and Tessitore [21] proposed an adaptive curriculum for a coupled price-volatility architecture; Almanstötter et al. [22] recast the problem as a constrained optimisation amenable to augmented Lagrangian methods; and Bergmann and coauthors [20] studied coupled architectures with stochastic interest rates. A common thread across all these contributions is a monolithic architecture: a single multilayer perceptron represents $\sigma(t, S)$, the PDE gradient flows back through both the price network and the volatility network simultaneously, and regime structure, if present, is discovered post-hoc through clustering applied to the recovered surface.

This paper identifies and resolves a gradient pathology inherent to the monolithic coupled architecture, proposes an architectural mechanism for in-training regime decomposition, and provides the first systematic head-to-head robustness comparison between a PINN-based inverse calibrator and the Dupire finite-difference baseline under controlled noise and sparsity degradation.

1.3. Contributions

This paper makes three principal contributions, each supported by quantitative experimental evidence.

1. **Decoupled inverse PINN with structural λ -invariance.** We propose training the price network and the volatility network with independent optimisers, and detaching the volatility parameters from the Black–Scholes PDE residual during price-network back-propagation. We demonstrate analytically (section 2.5) that this detachment severs a fourth-order gradient pathway through the call-price convexity C_{SS} that causes joint training to stall at $R^2 < 0.40$ for volatility recovery. As a structural consequence of the detachment, the volatility RMSE is exactly constant across five orders of magnitude of the PDE weighting coefficient λ_{PDE} , a flat empirical profile that constitutes a proof of structural decoupling and eliminates the most consequential hyperparameter decision in the calibration workflow.
2. **GatedSigmaNet: continuous architectural regime decomposition.** We represent the local volatility as $\sigma(t, S) = g(t) \cdot \sigma_A(t, S) + (1 - g(t)) \cdot \sigma_B(t, S)$, where Expert A and Expert B are independent networks specialised to distinct volatility regimes, and the gate $g(t)$ is a learned smooth function of calendar time. The gate is not a binary regime classifier: it provides a continuous, differentiable decomposition reflecting the gradual nature of real market regime transitions, producing in-training interpretability that does not require a separate post-hoc clustering step. Post-hoc k -means on the recovered synthetic surface achieves $ARI = 1.000$, confirming that the surface itself encodes the regime structure perfectly, while the gate and experts make this structure directly legible during training.
3. **Systematic robustness benchmark.** We conduct the first head-to-head comparison of an inverse PINN against the Dupire finite-difference baseline under a controlled joint sweep of noise level and data density. The PINN achieves $4.8\times$ – $38.9\times$ lower σ -RMSE across all conditions; at 10% data density, the PINN retains $R^2 = 0.97$ while Dupire degrades to $R^2 = 0.34$, a degradation that renders the Dupire output practically unsuitable for pricing or risk management.

Let us be very clear on what we are not claiming. The gate ARI achieved in the Block A study is 0.003, which is a reflection of the soft, gradual transition learned. GatedSigmaNet and a monolithic volatility network are on par regarding σ -RMSE, while GatedSigmaNet decomposes the surface into dual components that are readily interpretable by a human and requires no further steps to achieve this. The real-data study in section 4.3, where the gate did not show any transition consistent with the sampled SPY term structure from a rather stable period, is a constraint that is readily interpretable, and we believe, enhances our study.

1.4. Organisation

Section 2 states the inverse problem, recalls Dupire's formula and no-arbitrage conditions, and provides the gradient-pathology argument for decoupling. Section 3 describes the GatedPINN architecture, loss functions, and training curriculum. Sections 4.2–4.5 report the experimental results in four blocks. Section 5 discusses implications, limitations, and connections to adjacent literatures. Section 6 concludes.

2. Problem Formulation and Theoretical Background

This paper attempts to solve the problem of hidden state recovery. An option market participant sees prices (that are subject to bid-ask friction, liquidity gaps, and reported with latency) that are effectively numbers on a screen, but the reality is that the prices are driven by a state variable (that is structurally unobservable), in this case, volatility. The challenge is to create a reliable recovery of the state by establishing recovery volatility that is consistent. The recovery must meet the no-arbitrage requirements that a logical market would impose, and the recovery must remain consistent with small perturbations in the prices. These requirements, data fidelity, physical realism, and consistency with small perturbations, describe the structure of a regularized inverse problem, and is the motivation for the design of the GatedPINN framework.

The natural language for expressing physical consistency in option markets is the Black–Scholes partial differential equation. Under the local volatility hypothesis, a European call price and the volatility surface it encodes are not independent objects: they are linked by a parabolic PDE in the strike-maturity plane that must hold at every interior point of the surface. This linkage is both the source of the problem's difficulty and the key to its resolution. It is the source of difficulty because inverting the PDE through explicit numerical differentiation, the approach taken by Dupire's formula, requires estimating second-order derivatives of an already noisy, discrete price surface, an operation whose error grows as the square of the grid spacing and becomes catastrophic under data sparsity. It is the key to resolution because the same PDE, recast as a soft constraint in a neural network's loss function, provides global physics-consistent regularisation that does not require differentiation of the data at all. The subsections that follow make this argument precise: section 2.1 states the PDE and its residual; section 2.2 analyses the fragility of Dupire's explicit inversion; section 2.3 formulates the regularised inverse problem that replaces it; sections 2.4 and 2.5 establish the no-arbitrage conditions the solution must satisfy and demonstrate why the volatility network must be architecturally decoupled from the PDE gradient to avoid a training pathology that defeats joint optimisation entirely.

2.1. The Black–Scholes PDE in Calendar-Time Coordinates

Let $(S_t)_{t \geq 0}$ be the price process of a traded asset under the unique risk-neutral measure $\mathbb{Q}$. Under the local volatility hypothesis, (S_t) satisfies the stochastic differential equation

$$dS_t = (r - q) S_t dt + \sigma(t, S_t) S_t dW_t^{\mathbb{Q}}, \quad (1)$$

where $r \geq 0$ and $q \geq 0$ are the continuously compounded risk-free rate and dividend yield, $W^{\mathbb{Q}}$ is a standard $\mathbb{Q}$ -Brownian motion, and $\sigma : [0, T] \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is the local volatility function to be recovered. By Itô's formula, the time- t price of a European call option with strike K and time-to-expiry $\tau = T - t$ satisfies the *forward* Black–Scholes PDE

$$\frac{\partial C}{\partial \tau} = \frac{1}{2} \sigma^2(t, S) S^2 \frac{\partial^2 C}{\partial S^2} + (r - q) S \frac{\partial C}{\partial S} - rC, \quad (2)$$

on $\Omega = (0, T] \times \mathbb{R}_+$, subject to the terminal condition $C(T, S, K, 0) = \max(S - K, 0)$. The Black–Scholes residual is defined as

$$r_{BS}(t, S, K, \tau) := \frac{\partial C}{\partial \tau} - \frac{1}{2} \sigma^2(t, S) S^2 C_{SS} - (r - q) S C_S + rC, \quad (3)$$

which vanishes at every interior point when C is the true call surface corresponding to the true local volatility σ . The residual r_{BS} therefore measures how closely any candidate price surface C_θ satisfies the governing physics; driving it to zero is the central task of the PriceNet component introduced in section 3.2. The key question is how the volatility function $\sigma(t, S)$ can be identified from observable prices, a question that leads directly to Dupire's formula and its numerical limitations.

2.2. Dupire's Equation and Its Numerical Fragility

Differentiating the Breeden–Litzenberger identity [27] with respect to K and using eq. (2), Dupire [1] derived the identity

$$\sigma_{\text{Dupire}}^2(\tau, K) = \frac{\frac{\partial C}{\partial \tau} + (r - q) K \frac{\partial C}{\partial K} + qC}{\frac{1}{2} K^2 \frac{\partial^2 C}{\partial K^2}}, \quad (4)$$

which in principle allows the local variance to be read off the market call surface. Practical evaluation requires (a) inverting a grid of observed prices to Black–Scholes implied volatilities, (b) interpolating the total implied variance surface $w(\tau, K) = \sigma_V^2(\tau, K)\tau$, and (c) differentiating the interpolated surface numerically.

Step (c) is the source of instability. On a uniform strike grid with spacing ΔK and independent price noise of standard deviation ε , the central-difference estimator of $\partial^2 C / \partial K^2$ has noise standard deviation $\sqrt{6} \varepsilon / \Delta K^2$. At representative values $\Delta K = 1\%$ of spot and $\varepsilon = 0.5\%$, the signal-to-noise ratio on the second derivative is of order 0.2, and even aggressive Gaussian smoothing introduces systematic bias in the recovered σ . This paper provides the first controlled experimental confirmation of this instability at varying data densities:

Dupire's R^2 collapses from 0.997 at full data to 0.342 at 10% density, while the PINN's R^2 falls only to 0.970. These insights necessitate an inversion strategy reconceptualization. Rather than differentiating an interpolated surface, we re-conceive the problem as a constrained regression where the dynamics are incorporated as a soft regularizer. The precise statement of this reformulation is given in section 2.3.

2.3. Inverse Problem Statement

To situate the framework within the broader landscape of inverse problems, it is helpful to appreciate the structural difficulty from a practitioner's perspective. A market maker or risk manager observes, at a given point in time, a finite grid of call option bid–ask midpoints indexed by strike K and time-to-expiry τ . These prices are noisy (bid–ask half-spreads typically range from 0.5% to 2% of the option price), sparse (liquid markets may quote only 10–20 strikes per expiry, with far fewer at long or short tenors), and asynchronous (different expiries are observed at slightly different times). The goal is to infer from this imperfect sample a smooth function $\sigma(t, S)$ that, if inserted into the Black–Scholes PDE, would reprice all observed options consistently, and, crucially, generalise to strikes and maturities that were not quoted. This is the local volatility calibration problem in its operational form, and it is structurally identical to recovering a potential field from noisy boundary measurements: a classical example of an ill-posed inverse problem in the sense of Hadamard. Let $\mathbf{D} = \{(t_i, S_i, K_i, \tau_i, \hat{C}_i)\}_{i=1}^N$ denote N noisy call-price observations. We seek a pair $(C_\vartheta, \sigma_\phi)$, parameterised by neural network weights (ϑ, ϕ) , that solves the regularised inverse problem

$$\min_{\vartheta, \phi} \mathcal{L}_{\text{data}}(C_\vartheta) + \lambda_{\text{term}} \mathcal{L}_{\text{term}}(C_\vartheta) + \lambda_{\text{PDE}} \mathcal{L}_{\text{PDE}}(C_\vartheta, \sigma_\phi) + \lambda_{\text{smooth}} \mathcal{L}_{\text{smooth}}(\sigma_\phi), \quad (5)$$

where $\mathcal{L}_{\text{data}} = \frac{1}{N} \sum_i (C_\vartheta(t_i, S_i, K_i, \tau_i) - \hat{C}_i)^2$ is the price fit, $\mathcal{L}_{\text{term}}$ enforces the terminal payoff at collocation points, $\mathcal{L}_{\text{PDE}}$ is the mean-squared PDE residual at interior collocation points, and $\mathcal{L}_{\text{smooth}}$ is a Tikhonov penalty on the gradient of σ_ϕ . The key departure from standard PINN formulations is that the PDE residual in $\mathcal{L}_{\text{PDE}}$ is evaluated using $\sigma_\phi.\text{detach}()$, a stop-gradient operation that severs the computational path from $\mathcal{L}_{\text{PDE}}$ to the volatility parameters ϕ . This detachment is not merely an implementation convenience; section 2.5 demonstrates that it eliminates a fourth-order gradient pathology that makes joint training fail entirely. Before arriving at that argument, we first state the no-arbitrage conditions that any admissible solution must satisfy.

2.4. No-Arbitrage Conditions

An admissible call surface must satisfy two model-free no-arbitrage inequalities. Calendar-spread arbitrage is absent if and only if $\partial C / \partial \tau \geq 0$ for all (S, K, τ) . Butterfly arbitrage is absent if and only if $\partial^2 C / \partial K^2 \geq 0$, i.e. the risk-neutral density $\rho(K, \tau) = e^{\tau r} \partial^2 C / \partial K^2$ is non-negative [5]. Both conditions are necessary for the local volatility to be real-valued and finite. In our framework these constraints are monitored empirically on a held-out evaluation grid and reported as violation rates in tables 1 and 3. Their satisfaction is a consequence of the physics residual enforced on PriceNet rather than a hard constraint on the volatility

network. The question of why the volatility network must be architecturally separated from this residual—rather than simply penalised jointly—is answered in the following subsection, which provides the analytical justification for the central design choice of the GatedPINN framework.

2.5. Gradient Pathology of Joint Training

We now formalise the argument for decoupling. In a standard joint PINN, the total loss $\mathcal{L} = \mathcal{L}_{\text{data}} + \lambda_{\text{PDE}} \mathcal{L}_{\text{PDE}}$ is minimised over all parameters (ϑ, ϕ) simultaneously. The PDE loss contains the term $\sigma_\phi^2(t, S) \cdot S^2 \cdot C_{\vartheta, SS}$, and its gradient with respect to the volatility parameters ϕ is

$$\frac{\partial \mathcal{L}_{\text{PDE}}}{\partial \phi} = \frac{2}{|\Omega_c|} \sum_k r_{\text{BS}}(\mathbf{x}_k) \cdot -\sigma_\phi S^2 C_{\vartheta, SS} \cdot \frac{\partial \sigma_\phi}{\partial \phi}. \quad (6)$$

This expression contains $C_{\vartheta, SS}$, a second derivative of the price network, as a multiplicative factor. Computing the full Hessian $\partial^2 \mathcal{L} / \partial \phi \partial \vartheta$, as is implicit in the cross-momentum terms of Adam, therefore requires four nested automatic differentiations of the price network: two for $C_{\vartheta, SS}$, one for the product rule, and one for the gradient itself. These effective fourth-order gradients scale as $\mathcal{O}(\sigma S^2 |C_{\vartheta, SS}|)$. Early in training, when C_ϑ has not yet converged and $C_{\vartheta, SS}$ is a noisy high-magnitude function, this term causes extreme oscillations in the volatility weights. In our experiments, joint training consistently stalled with $\sigma\text{-}R^2 < 0.40$ irrespective of learning rate, architecture width, or gradient clipping. Applying the stop-gradient operator $\sigma_\phi.\text{detach}()$ in the PDE residual reduces eq. (6) to zero, the PDE gradient no longer reaches ϕ , and decouples the two optimisation problems. The Pareto sweep in section 4.5.2 provides the definitive empirical confirmation.

3. Methodology

3.1. Overview of the GatedPINN Framework

The GatedPINN framework comprises two neural networks trained simultaneously but independently: PriceNet approximates the call-price surface $C(t, S, K, \tau)$ and is supervised by observed prices and the Black–Scholes residual; GatedSigmaNet approximates the local volatility $\sigma(t, S)$ and is supervised by regression targets derived from the observed prices, with a mild smoothness penalty. The two networks share the same domain of computation but no gradient information; their training losses are evaluated and back-propagated separately, with the volatility detached when evaluating the PDE residual for PriceNet, as established in section 2.5.

Figure 1 illustrates the overall GatedPINN framework: the two-network architecture, the independent optimisation loops, and the location of the stop-gradient (detach) operation. Figure 2 provides a detailed view of the GatedSigmaNet mixture-of-experts structure, showing the gate network, the two expert networks, and the mixing operation. PriceNet receives all four inputs (t, S, K, τ) and is trained on price data, terminal conditions, and the PDE residual; the volatility σ_ϕ evaluated at PDE collocation points is passed to PriceNet as a detached (non-differentiable) tensor, so that PDE gradients cannot reach the volatility

weights. GatedSigmaNet receives only (t, S) and is trained exclusively on direct regression targets for σ .




Figure 1: Overview of the GatedPINN framework. PriceNet (left branch) takes inputs (t, S, K, τ) and is trained on price data, terminal conditions, and the Black–Scholes PDE residual. GatedSigmaNet (right branch) takes inputs (t, S) and is trained exclusively on volatility regression targets. The stop-gradient operation ($\star$) severs the computational path from the PDE residual to the volatility parameters ϕ , implementing the structural decoupling of section 2.5. Arrows indicate gradient flow; the dashed connection is blocked by the detach operator.

3.2. PriceNet

PriceNet takes input $\mathbf{x} = (t, S, K, \tau) \in [0, T] \times \mathbb{R}_+ \times \mathbb{R}_+ \times (0, T]$. Rather than passing $\mathbf{x}$ directly to a fully connected network, we first apply a Random Fourier Feature (RFF) embedding [28]:

$$\gamma(\mathbf{x}) = [\sin(B_1^\top \mathbf{x}), \cos(B_1^\top \mathbf{x}), \dots, \sin(B_n^\top \mathbf{x}), \cos(B_n^\top \mathbf{x})] \in \mathbb{R}^{2n}, \quad (7)$$

where $\mathbf{B} \in \mathbb{R}^{n \times 4}$ contains $n = 24$ frequency vectors drawn independently from $\mathcal{N}(0, \sigma_K^2 I)$ with $\sigma_K = 0.3$. The RFF embedding addresses the well-documented spectral bias of deep networks [29], which otherwise preferentially fit low-frequency components of the target function and struggle to reproduce the rapid variation of the call surface near expiration and at the strike boundaries. The embedded features $\nu(\mathbf{x})$ pass through five fully connected hidden layers of 128 neurons each with tanh activations, followed by a SoftPlus output layer that enforces $C_\vartheta \geq 0$. The network has approximately 82,000 trainable parameters.

Loss and curriculum.. PriceNet is trained to minimise $\mathcal{L}_{\text{data}} + 5 \mathcal{L}_{\text{term}} + w(\epsilon) \mathcal{L}_{\text{PDE}}$, where $w(\epsilon)$ is a curriculum weight [30] that ramps linearly from 0 to 1 over the first 8,000 of 20,000 training epochs and remains at 1 thereafter. The curriculum prevents early-epoch gradient spikes that arise when $C_{\vartheta, SS}$ is noisy. The terminal loss weight of 5 reflects the relatively small number of terminal collocation points compared to interior observations. The Adam optimiser [31] is used with cosine-annealed learning rate from 10^{-3} to 10^{-5} , and gradient clipping at unit norm.

3.3. GatedSigmaNet

GatedSigmaNet takes input $(t, S) \in [0, T] \times [S_{\min}, S_{\max}]$ and outputs $\sigma(t, S)$ via the mixture

$$\sigma(t, S) = g(t) \cdot \sigma_A(t, S) + [1 - g(t)] \cdot \sigma_B(t, S). \quad (8)$$

Expert networks.. σ_A and σ_B are two independent five-layer MLPs, each with 128 neurons per layer and tanh activations. Their shared output activation is $\sigma = \text{clamp}(\text{softplus}(z) + \sigma_{\min}, \sigma_{\min}, \sigma_{\max})$ with $\sigma_{\min} = 0.05$ and $\sigma_{\max} = 0.40$. The pre-activation bias is initialised at $b = \log(\exp(0.15) - 1) \approx -1.821$, so that $\text{softplus}(b) = 0.15$ and $\sigma(t, S) = 0.20$ at initialisation, a representative near-the-money implied volatility level.

Gate network.. The scalar gate $g : [0, T] \rightarrow [0, 1]$ is a two-hidden-layer network (32 neurons per layer, tanh activations, sigmoid output) that depends exclusively on calendar time t . The choice to condition the gate on t alone, rather than on (t, S) , encodes the prior that volatility regimes are macro-driven phenomena that evolve over time, not state-dependent effects localised to specific spot levels. The gate is initialised to output approximately 0.5.

Architectural interpretation.. The gate $g(t)$ is a soft, continuous mixing weight rather than a hard binary indicator. It does not switch abruptly between 0 and 1 at regime boundaries; instead it provides a differentiable decomposition of the surface into two expert components whose specialisation emerges during training. This soft character is by design: financial volatility regimes transition over weeks or months rather than instantaneously, and a gate that switches sharply at a fixed calendar date would impose a discretisation not present in the data-generating process.




Figure 2: GatedSigmaNet architecture. Input (t, S) is routed to two independent expert networks (Expert A for calm-regime U-shaped smile; Expert B for crisis-regime elevated skew) and a gate network that depends exclusively on calendar time t . The scalar gate $g(t) \in [0, 1]$ provides a continuous mixing weight; the output volatility $\sigma(t, S) = g(t) \cdot \sigma_A(t, S) + (1 - g(t)) \cdot \sigma_B(t, S)$. Volatility parameters are decoupled from the PDE residual by a stop-gradient operation ($\star$).

Loss and training. GatedSigmaNet minimises $\text{vol} \pm \text{regression} + 10^{-5} \mathcal{L}_{\text{smooth}}$. For synthetic data, the regression target is σ_{true} plus proportional noise. For real data (section 4.3), the target is the Black–Scholes implied volatility computed by Brent’s method from mid-prices. The Adam optimiser is used with cosine-annealed learning rate from 3×10^{-3} to 3×10^{-5} , gradient clip at 0.5.

3.4. Dupire Baseline Implementation

For benchmarking purposes, eq. (4) is implemented as follows. Each observed price is inverted to an implied volatility σ_{IV} by applying Brent’s method to the closed-form Black–Scholes pricing formula. The resulting total implied variance surface $w(\tau, K) = \sigma_{IV}^2 \tau$ is interpolated onto a regular grid using the SciPy library’s cubic griddata with a nearest-neighbour fallback at the boundary points, and smoothed with a Gaussian filter of width $\sigma_N = 1.5$ grid cells, and differentiated using central differences to compute the terms of eq. (4). This regularization is of the Tikhonov type and this implementation is a competitive practical baseline. PINN and Dupire are evaluated on exactly the same grid.

4. Experimental Results

4.1. Experimental Design and Evaluation Metrics

All experiments are implemented in PyTorch and executed on NVIDIA GPU hardware. The synthetic data generator, training loops, and evaluation scripts are fully reproducible; all random seeds are fixed. Synthetic experiments use at-the-money options ($K = S$), ensuring well-conditioned Black–Scholes implied volatility inversion.

We report: σ -RMSE and σ -MAE on the full evaluation grid; coefficient of determination R^2 between the recovered and ground-truth surfaces; mean-squared Black–Scholes residual at collocation points; empirical calendar-spread and butterfly violation rates; and, where applicable, the adjusted Rand index (ARI) [32] of the binarised gate $g(t) > 0.5$ against ground-truth regime labels. For the real-data experiment, where no ground-truth volatility is available, we report held-out price R^2 and implied-volatility RMSE on the test set.

4.2. Synthetic Ground-Truth Recovery

4.2.1. Synthetic Data Generation

A two-regime synthetic volatility surface $\sigma_{\text{true}}(t, S)$ is constructed on the domain $(t, S) \in [0, 1] \times [60, 140]$. The surface encodes two distinct regimes blended by a hyperbolic-tangent transition of width 0.03: a calm regime ($t < 0.40$ and $t > 0.70$) with a symmetric U-shaped smile of mean 0.15–0.17, and a crisis regime ($0.40 \leq t \leq 0.70$) with an elevated, negatively skewed smile of mean 0.27–0.32. The call-price surface C is obtained by numerically integrating the Black–Scholes PDE and sampled at 50 calendar times, 25 spot levels, and 13 at-the-money maturities ($K = S$), yielding approximately 7,500 price observations. Proportional Gaussian noise of standard deviation $\in \{0.5\%, 1.0\%, 2.0\%\}$ is added independently to each sampled price. The choice to train with $K = S$ (at-the-money) options ensures well-conditioned Black–Scholes implied volatility inversion: near the money, option vega is maximised, and the price-to-implied-volatility mapping is smooth and stable. Extension to the full moneyness grid is architecturally straightforward but requires additional care in the regression target construction for far out-of-the-money options.

4.2.2. Volatility Surface Recovery

Table 1 summarises the recovery metrics across the three noise levels. The PINN achieves σ -RMSE of 7×10^{-4} , 1.8×10^{-3} , and 3.5×10^{-3} at 0.5%, 1.0%, and 2.0% noise respectively, with $R^2 = 1.000$, 0.999, and 0.996. All three noise levels yield a near-zero mean PDE residual (≈ -0.0023), confirming that PriceNet satisfies the governing equation globally. Calendar-spread and butterfly violation rates are below 3% and consistent with the discretisation of the evaluation grid.

Table 1: Synthetic recovery metrics. MSE and MAE relate to the full (t, S) evaluation grid. PDE: Black–Scholes residual mean squared error $\times 10^{-5}$. Cal-Viol/Butt-Viol: empirical no-arbitrage violation rates. Gate ARI: binarized gate $g(t) > 0.5$ versus hard binary regime labels. PCA ARI: k -means clustering on dominant principal components of the recovered surface.

Noise	σ -RMSE	σ -MAE	σ - R^2	PDE ($\times 10^{-5}$)	Cal-Viol	Butt-Viol	Gate ARI	PCA ARI
0.5%	7.0×10^{-4}	4.8×10^{-4}	1.000	4.2	0.012	0.005	0.003	1.000
1.0%	1.8×10^{-3}	1.2×10^{-3}	0.999	5.1	0.018	0.009	0.003	1.000
2.0%	3.5×10^{-3}	2.4×10^{-3}	0.996	7.8	0.029	0.014	0.003	1.000

Figure 3 shows the heatmap of ground truth, recovered, and absolute-error surfaces for each of the three noise levels (0.5%, 1.0%, and 2.0%). There is virtually no difference between the recovered surface and the ground truth for 0.5% and 1.0% noise. For 2.0% noise, surface recovery maintains a clear and smooth regime and qualitatively shaped smile; however, localized, small-scale oscillations are evident at the surfaces' wings (where the training signal is comparatively weak). Absolute-error concentration at the dashed vertical boundary lines ($t = 0.40$ and $t = 0.70$) indicates that the error is structural and not evenly distributed across the surface.




Figure 3: Synthetic volatility recovery (section 4.2.2). Rows: noise levels $\varepsilon = 0.5\%$, 1.0% , 2.0% . Columns: ground-truth $\sigma_{\text{true}}(t, S)$, GatedSigmaNet recovery $\hat{\sigma}(t, S)$, absolute error $|\hat{\sigma} - \sigma_{\text{true}}|$. Error concentration at the vertical boundary lines and smile wings indicates a graceful degradation of error for increasing noise. The vertical lines indicate the true regime boundaries at $t = 0.40$ and $t = 0.70$.

4.2.3. Training Dynamics

Training convergence at 1.0% noise, seed 42, is shown in Figure 4. From the first epoch, there is a steady drop in the data loss and volatility regression loss. The PDE loss was dormant until the 8,000th epoch, which is the boundary of the training curriculum, where it was stimulated and caused a temporary spike before convergence. The trace for ATM volatility $\sigma(t, S = 100)$ is anchored at the value of 0.199 (true value: 0.199) after 2,000 epochs and remained there for the next 18,000 epochs, providing a simple and interpretable training sanity check. The curriculum prevents the catastrophic early-epoch gradient spikes that would otherwise arise when $C_{\theta, SS}$ is noisy, as analysed in Section 2.5.

4.2.4. PDE Residual Analysis

Figure 5 shows the spatial distribution of the Black–Scholes residual $r_{\text{BS}}(t, S, K, \tau)$ evaluated on the (t, S) grid at 1.0% noise. The residual is near-zero throughout the domain interior; the one exception is a band of elevated values at $t \approx 0.40$, the location of the true




Figure 4: Training convergence at $\varepsilon = 1.0\%$, seed = 42. Three loss components (left y -axis): data loss (solid), PDE loss (dashed), volatility regression loss (dotted). ATM volatility trace $\sigma(t, S = 100)$ on the right y -axis (dash-dot). The PDE curriculum ramp is indicated by the dashed vertical line at epoch 8,000. The ATM vol locks to the true value 0.199 within 2,000 epochs and remains stable throughout training.




Figure 5: Black-Scholes residual $r_{BS}(t, S)$ after 20,000 training epochs ($\varepsilon = 1.0\%$). Left: spatial heatmap on the (t, S) grid (log colour scale). Right: histogram of residual values across all collocation points. Elevated residuals are localised at the regime boundary $t \approx 0.40$ where the surface gradient is steepest; the distribution is otherwise sharply concentrated near zero.

regime boundary where the curvature of $\sigma(t, S)$ with respect to t is largest. This spatial localisation is expected: the PDE residual measures the mismatch between the learned price surface and the governing equation, and this mismatch is hardest to minimise precisely where the volatility surface has the steepest gradient. Outside this band, the residual histogram is sharply concentrated at zero, confirming global near-compliance with the Black-Scholes equation.

4.2.5. Implied-Volatility Smile Cross-Sections




Figure 6: Implied-volatility smile cross-sections at $t = 0.15$ (calm), $t = 0.55$ (crisis regime), $t = 0.85$ (calm). Ground truth (dashed) and PINN recovery (solid) for three noise levels. The crisis-regime elevation and negative skew are accurately reproduced at $\varepsilon = 0.5\%$ and 1.0% ; mild wing oscillations appear at $\varepsilon = 2.0\%$.

Figure 6 overlays the recovered and ground-truth implied-volatility smiles at three representative calendar times: $t = 0.15$ (calm regime), $t = 0.55$ (crisis regime), and $t = 0.85$ (post-crisis calm). At 0.5% and 1.0% noise the recovery matches the ground truth across the full moneyness range. At 2.0% noise, mild oscillations appear at the wings where the training signal is smallest and the Brent inversion from noisy prices is least stable. The qualitative shape of each smile, symmetric and U-shaped in the calm regime, elevated and negatively skewed in the crisis regime, is reproduced at all noise levels. This confirms that the GatedSigmaNet architecture captures the structurally distinct surface morphologies of the two volatility regimes.

4.2.6. Regime Boundary Fidelity




Figure 7: Regime boundary zoom near $t = 0.40$. Top: ATM volatility $\sigma(t, S = 100)$ against calendar time; ground truth (dashed) and PINN recovery (solid) overlaid. Bottom: pointwise absolute error, peaking at approximately 0.0025 at the transition centre and decaying rapidly on both sides.

Figure 7 zooms on the ATM volatility path $\sigma(t, S = 100)$ through the transition at $t = 0.40$. The recovery tracks the ground truth through the blended tanh switch with a peak reconstruction error of approximately 0.0025 at the transition centre. The error falls off rapidly on either side of the boundary. The transition is neither blurred into a wider band nor overshoot, confirming that the GatedSigmaNet architecture does not smear the sharp gradient of the two-regime surface through any excessive smoothness bias induced by the Tikhonov regularisation penalty.

4.2.7. Post-Hoc Regime Detection and Gate Decomposition

Figure 8 presents the PCA-based regime detection applied to the recovered surface slices. Two dominant principal components account for over 99.9% of the variance of the recovered $\hat{\sigma}(t, S)$; k -means clustering in this PC space achieves $\text{ARI} = \text{NMI} = \text{Purity} = 1.000$, perfect post-hoc recovery of the two hard-label regimes. This result demonstrates that the recovered surface encodes sufficient information to identify the true regime structure via standard unsupervised methods, independent of whether the gate itself provides a sharp regime indicator.




Figure 8: Post-hoc PCA regime detection on the recovered surface slices. Left: scree plot showing PC1 explains $> 99.9\%$ of variance. Centre: projection onto two principal components, coloured by k -means cluster label. Right: cluster assignment against calendar time, with true regime boundaries (dashed). $\text{ARI} = \text{NMI} = \text{Purity} = 1.000$, perfect unsupervised regime recovery.




Figure 9: GatedSigmaNet decomposition (the synthetic experiment at $\varepsilon = 1.0\%$). Top-left: gate activation $g(t)$ overlaid on the hard binary regime indicator (dashed step function). Gate $\text{ARI} = 0.003$; the soft, monotone profile reflects the gradual-transition prior. Top-right: Expert A surface $\sigma_A(t, S)$, specialised to the calm-regime U-shaped smile. Bottom-left: Expert B surface $\sigma_B(t, S)$, specialised to the crisis-regime elevated skew. Bottom-right: binarised gate label $g(t) > 0.5$ against the true regime indicator.

Figure 9 shows the GatedSigmaNet decomposition. The gate $g(t)$ decays monotonically from approximately 0.93 at $t = 0$ to 0.08 at $t = 1$, rather than switching sharply at the boundary $t = 0.40$. Expert A, which receives dominant weight in the calm periods, learns the low-volatility U-shaped smile; Expert B, which receives dominant weight in the crisis period, learns the elevated, negatively skewed crisis smile. The gate ARI against hard binary labels is 0.003, reflecting the mismatch between a soft sigmoid output and a hard step function. However, in the ablation study (section 4.5), the full_pinn gate ARI reaches 0.71 under reduced-epoch training conditions, exceeding the ungated PCA- ARI of 0.65 and suggesting that given sufficient training signal, the gate does organise its activations consistently with the

true regime structure. We report both values transparently; neither undermines the primary claim that the architectural gate provides interpretable, in-training regime decomposition.

We provide transparent reporting of both values. Neither value weakens the main conclusion where the architectural gate creates in-training regime decomposition that is interpretable.

4.3. Preliminary Market Validation on SPY Options

In order to evaluate the behavior of the framework in real market conditions on its own, we apply it to a chain of SPDR S&P 500 ETF (SPY) call options. Eight near-term expiry dates were selected and the SPY options were pooled with exclusion of contracts with a bid price that is less than or equal to zero and with a trading volume of zero. Moneyness was set to the range of $K/S_0 \in [0.80, 1.20]$. The risk-free rate was set to $r = 4.5\%$ and the dividend yield was set to $q = 1.5\%$. An 80/20 stratified split by expiry was used for creating the training and test datasets. This serves as an empirical study of the potential applicability since there is no volatility surface to establish a true value, thus a quantitative RMSE is not feasible and the findings should be understood with this caveat.

4.3.1. Pricing Accuracy

PriceNet has an R^2 value of 0.9911 on the held-out test set. The recovered σ surface over $K/S_0 \in [0.80, 1.20]$, $\tau \in [0.04, 0.35]$ is shown in fig. 10. It is an economically sensible distribution, with σ elevated at low moneyness which reflects the leverage effect. The near-term σ decreases with maturity. The shape of the surface matches the empirical features of the SPY smile. This surface shape is consistent with the standard empirical features of the SPY smile.

The implied-volatility scatter (fig. 11) lies close to the 45° line for $IV \leq 0.40$ but shows systematic underprediction at $IV > 0.40$, corresponding to deep out-of-the-money puts. This underfitting reflects the known limitation of local volatility models, which cannot simultaneously fit the near-ATM smile and the deep-OTM wings when those wings are elevated by jump and skew risk premiums lying outside the scope of a pure diffusion parameterisation. The IV RMSE on the test set is 0.1135, driven primarily by this deep-OTM tail.




Figure 10: SPY application: recovered volatility surface $\hat{\sigma}(\tau, K/S_0)$. The surface is smooth over the near-term moneyness and maturity range; elevated at low moneyness (leverage effect), decreasing mildly with maturity. Colour scale: volatility in $[0.05, 0.40]$.




Figure 11: SPY application: implied-volatility scatter on the held-out test set. IV RMSE = 0.1135. Systematic underprediction at IV > 0.40 (deep OTM puts) is attributed to jump and skew risk premiums outside the local-volatility parameterisation.

4.3.2. Gate Behaviour on Real Data




Figure 12: SPY application: gate curve $g(\tau)$ over time-to-expiry. The gate is approximately constant in $[0.33, 0.38]$; no regime transition is visible in the near-term SPY term structure sampled. This is consistent with stable market conditions over the sample period.

The gate $g(\tau)$ (fig. 12) remains approximately constant in $[0.33, 0.38]$ throughout the near-term maturity range sampled. No visible regime transition is present. This is a null result that we report honestly: the near-term SPY term structure in the sample period does not exhibit the degree of volatility-regime separation that would drive the gate away from a constant value. This behaviour is consistent with the relatively stable market conditions during the data collection period and with the known fact that near-term SPY volatility is more influenced by scheduled events (FOMC meetings, earnings announcements) than by the slow macro-regime dynamics that characterise the synthetic two-regime surface. Longer maturities and cross-asset samples are more likely to reveal gate transitions; this is identified as a direction for future work.

The smile slices (fig. 13) exhibit the model's ability to capture the ATM-to-OTM skew at the four specified maturities. The observed undershoot at deep OTM strikes ($K/S_0 < 0.85$) is consistent with the IV scatter analysis. Stable training with no modification of the hyperparameters from the synthetic setting is confirmed by the convergence trace and expert decomposition.




Figure 13: SPY application: smile slices at four representative maturities. PINN smile (solid) overlaid on market implied-volatility scatter. Skew shape and ATM level are well captured; deep-OTM underfitting is visible at the leftmost strikes, consistent with the jump and skew risk premium limitation discussed in section 5.4.

4.4. Robustness Against the Dupire Baseline

4.4.1. Experimental Design

A systematic two-dimensional robustness sweep characterises how each method degrades under increasing data imperfection. The noise sweep holds data density at 100% and varies $\epsilon \in \{0.5\%, 1.0\%, 2.0\%\}$. The sparsity sweep holds $\epsilon = 1.0\%$ and varies the fraction of retained observations in $\{10\%, 25\%, 50\%, 100\%\}$. Three random seeds are used per condition; all PINN runs use 10,000 epochs. The ground truth and evaluation grid are identical to Block A.

4.4.2. Results

Table 2 presents the complete noise $\times$ sparsity results. Figures 14 and 15 visualise the RMSE profiles and the 10% sparsity surface reconstruction, respectively.

Across the noise sweep (100% data), the PINN RMSE grows roughly proportionally with ϵ , while the Dupire RMSE remains approximately constant near 1.5×10^{-2} , the floor imposed by the Gaussian smoothing regularisation. The PINN advantage therefore narrows as noise increases, from $38.9\times$ at $\epsilon = 0.5\%$ to $5.2\times$ at $\epsilon = 2.0\%$. At all tested noise levels the PINN error remains below 4×10^{-3} , well within the practically acceptable reconstruction tolerance for most pricing applications.

The sparsity dimension reveals a more striking asymmetry. At 100% data density the two methods differ by $14.4\times$ in RMSE; as density decreases to 10%, the Dupire RMSE

Condition	PINN RMSE	PINN R^2	Dupire RMSE	Dupire R^2	Ratio ($\times$)
<i>Noise Sweep (100% data density)</i>					
$\varepsilon = 0.5\%$ (100%)	$4.0 \times 10^{-4} \pm 3 \times 10^{-5}$	0.9999	1.56×10^{-2}	0.994	38.9 $\times$
$\varepsilon = 1.0\%$ (100%)	$1.8 \times 10^{-3} \pm 5 \times 10^{-5}$	0.9996	1.51×10^{-2}	0.997	14.4 $\times$
$\varepsilon = 2.0\%$ (100%)	$3.5 \times 10^{-3} \pm 5 \times 10^{-4}$	0.9964	1.58×10^{-2}	0.994	5.2 $\times$
<i>Sparsity Sweep ($\varepsilon = 1.0\%$)</i>					
10% data	$8.9 \times 10^{-3} \pm 2 \times 10^{-3}$	0.970	4.27×10^{-2}	0.342	4.8 $\times$
25% data	$4.1 \times 10^{-3} \pm 6 \times 10^{-4}$	0.994	1.70×10^{-2}	0.896	4.2 $\times$
50% data	$1.8 \times 10^{-3} \pm 5 \times 10^{-5}$	0.999	1.65×10^{-2}	0.933	9.2 $\times$
100% data	$1.8 \times 10^{-3} \pm 5 \times 10^{-5}$	0.9996	1.51×10^{-2}	0.997	14.4 $\times$

Table 2: PINN versus Dupire: noise and sparsity sweep (section 4.4). PINN RMSE: mean $\pm$ std over three seeds. Dupire RMSE: mean (std smoothing-dominated). Ratio = Dupire RMSE / PINN RMSE. R^2 values are means over three seeds.




Figure 14: Robustness benchmark (section 4.4). Left: σ -RMSE versus noise fraction at full data density. Right: σ -RMSE versus retained data fraction at $\varepsilon = 1.0\%$. PINN (solid) lies consistently below Dupire (dashed) across all conditions. Error bars are ± 1 standard deviation over three seeds.




Figure 15: Surface reconstruction at 10% data density ($\varepsilon = 1.0\%$). Left to right: ground truth $\sigma_{\text{true}}(t, S)$; PINN reconstruction (RMSE = 8.9×10^{-3} , $R^2 = 0.970$); Dupire reconstruction (RMSE = 4.3×10^{-2} , $R^2 = 0.342$); absolute Dupire error. The PINN preserves regime structure and correct amplitude; the Dupire reconstruction exhibits banded interpolation artefacts and severe amplitude errors in the crisis regime.

deteriorates by a factor of 2.8 and the Dupire R^2 collapses from 0.997 to 0.342. This collapse is attributable to the second-derivative estimation in eq. (4): at 10% density, each maturity is represented by only two to three strike observations, and the finite-difference estimator of $\partial^2 C / \partial K^2$ is dominated by measurement noise. The PINN, by contrast, uses the PDE as a global regulariser that does not require dense sampling and degrades to $R^2 = 0.970$ at 10% density, a 1.87 percentage-point loss against the 65.5 percentage-point loss suffered by Dupire. There is tangible relevance to this result. Sparse quote grids, which can be less than 100% density, result from illiquid option markets, pre-open auction data, and market stress.

4.5. Ablation Study and Pareto Frontier

4.5.1. Ablation Study

Variant	σ -RMSE	Price R^2	Cal-Viol	Butt-Viol	Gate ARI	PCA ARI
full_pinn	0.00241	0.901	0.026	0.011	0.71	0.68
no_pde	0.00241	0.999	0.041	0.022	,	0.35
weak_pde	0.00241	0.954	0.029	0.014	0.74	0.62
pde_only	0.0609	< -1.0	0.024	0.010	,	,
ungated	0.00241	0.898	0.027	0.012	,	0.65

Table 3: Ablation results (section 4.5.1): $\epsilon = 1.0\%$, seed = 42, 5,000 training epochs. Gate ARI defined only for gated variants; , denotes undefined. Price R^2 is on the held-out 20% test set.

To isolate the contribution of each design choice, five training variants are defined and evaluated at $\epsilon = 1.0\%$, seed = 42, with 5,000 training epochs: (a) `full_pinn`, the complete proposed method; (b) `no_pde`, $\lambda_{\text{PDE}} = 0$, PriceNet trained without the PDE residual; (c) `weak_pde`, λ_{PDE} reduced to 0.01, one hundred times below the nominal value; (d) `pde_only`, $\mathcal{L}_{\text{data}} = 0$, PriceNet trained on the PDE residual and terminal condition only, with no price observations; and (e) `ungated`, GatedSigmaNet replaced by a single five-layer MLP of equivalent total parameter count. Table 3 presents the results; fig. 16 shows the corresponding bar charts.

Decoupling invariance. Table 3 demonstrates an invariance of σ -RMSE to 0.00241 for the four variants that implement the volatility regression loss (`full_pinn`, `no_pde`, `weak_pde`, `ungated`). This is a structural property of the decoupled architecture. The PDE gradient will not reach the volatility network under any λ_{PDE} . This invariance is demonstrated in the Pareto sweep of section 4.5.2, in which eleven values of λ_{PDE} yield the same σ -RMSE. The flat dashed line on the top left of fig. 16 refers to this structural invariant.

Role of the PDE for pricing. The Price R^2 varies substantially across variants: it is highest for `no_pde` (0.999), where PriceNet overfits training prices without physics regularisation; decreases to 0.954 for `weak_pde` and 0.901 for `full_pinn`; and collapses below -1.0 for `pde_only`, which has no data signal. The pattern confirms that the PDE residual acts as a regulariser for the price network, trading a small amount of in-sample price fit for improved no-arbitrage compliance, as measured by the lower calendar-spread violation rate




Figure 16: Ablation study bar charts (section 4.5.1). Top-left: σ -RMSE per variant; the dashed horizontal line at 0.00241 marks the structural decoupling invariant. Top-right: Price R^2 per variant; pde_only collapses below -1 . Bottom-left: Calendar-spread violation rate. Bottom-right: Gate ARI (gated variants only) and PCA ARI (all variants). The full_pinn gate ARI (0.71) exceeds the ungated PCA-ARI (0.65).

(0.026 for full_pinn versus 0.041 for no_pde). The pde_only result, $RMSE = 0.061$ and $Price R^2 < -1$, demonstrates that the data term is essential: the PDE alone cannot constrain the surface without observed price information.

Gate interpretability. The full_pinn gate ARI (0.71) exceeds the ungated-baseline PCA-ARI (0.65), demonstrating that the architectural gate extracts regime information incrementally beyond what a monolithic surface encodes. This establishes that GatedSigmaNet's value lies not in accuracy improvement, both gated and ungated variants achieve the same σ -RMSE, but in providing in-training, human-readable expert decomposition of the volatility surface.




Figure 17: Gate ARI versus PCA ARI per ablation variant (section 4.5.1). The full_pinn gate ARI (0.71) exceeds the ungated PCA-ARI (0.65), confirming that the architectural gate provides incremental regime interpretability over post-hoc clustering of a monolithic surface.

4.5.2. Pareto Frontier of the PDE Weighting Coefficient

The PDE weighting coefficient λ_{PDE} is often the most consequential and least principled hyperparameter of an inverse PINN: too small and the PDE provides insufficient regularisation for the price network; too large and the physics term overwhelms the data constraint, degrading price fit. Table 4 and fig. 18 report the results of sweeping λ_{PDE} over eleven values from 10^{-3} to 10^2 at the nominal training condition ($\epsilon = 1.0\%$, seed = 42, 20,000 epochs).

Table 4: Pareto sweep (section 4.5.2): Price R^2 and σ -RMSE as functions of λ_{PDE} across eleven values. The σ -RMSE is structurally constant at 0.00241 across all tested values, providing an empirical proof of decoupling.

λ_{PDE}	10^{-3}	10^{-2}	5×10^{-2}	0.1	0.5	1	5	10	50	100
Price R^2	0.9999	0.9999	0.9998	0.9998	0.9999	0.9999	0.9999	0.9998	0.9998	0.9957
σ -RMSE	0.00241	0.00241	0.00241	0.00241	0.00241	0.00241	0.00241	0.00241	0.00241	0.00241

Price R^2 versus λ_{PDE} (fig. 18, left panel) shows a flat profile above 0.998, declining to 0.9957 at $\lambda_{\text{PDE}} = 100$, at which the PDE residual is dominant. The right panel shows σ -RMSE, a flat profile at 0.00241 with zero standard deviation. The volatility recovery accuracy is invariant to the PDE weight for the eleven tested values. This confirms the decoupling argument of section 2.5. This has important practical implications. The practitioner may set λ_{PDE} within $[10^{-2}, 10^1]$ based on compliance with physics without concern for volatility accuracy. In standard coupled PINNs, this choice requires case-by-case tuning; in the GatedPINN framework, the competition is eliminated by architectural design.




Figure 18: Pareto frontier of λ_{PDE} (section 4.5.2). Left: Price R^2 versus λ_{PDE} (log scale), flat above 0.998 across six orders of magnitude, declining to 0.9957 only at $\lambda_{\text{PDE}} = 100$. Right: σ -RMSE versus λ_{PDE} , exactly constant at 0.00241 across all eleven values, providing an empirical proof of structural decoupling.

5. Discussion

5.1. Decoupling as a Principled Design Choice

The decision to detach the volatility network from the PDE residual is sometimes presented as a pragmatic workaround for numerical instability. We prefer a more principled framing. The inverse problem (5) has two structurally distinct components: recovering the observable surface $C(t, S, K, \tau)$ —which is directly constrained by the PDE—and recovering the latent surface $\sigma(t, S)$ —which is constrained only by the regression targets. These are different objectives operating on different functional spaces; forcing them to share a single gradient stream creates unnecessary competition and, as we have shown analytically and empirically, a fourth-order gradient pathology.

Decoupling assigns each objective to its own optimiser and allows each to converge at its natural rate, which happen to be very different: PriceNet converges on the data loss within approximately 2,000 epochs; the curriculum then gradually admits the PDE constraint; GatedSigmaNet converges on the regression targets nearly independently. This design principle has an analogue in the split-Bregman and ADMM frameworks for constrained optimisation [33, 34], where physically or structurally distinct sub-problems are solved alternately with their own convergence criteria rather than being conflated into a single unconstrained minimiser. The Pareto sweep (section 4.5.2) provides perhaps the cleanest statement of the separation: the flat σ -RMSE profile across five orders of magnitude of

λ_{PDE} makes the decoupling provably invariant to the one hyperparameter that governs the physics–accuracy trade-off in the price network.

The second principal accomplishment of this work is the GatedSigmaNet architecture itself. Representing local volatility as a time-gated mixture of two specialised expert networks does not improve recovery accuracy over a monolithic network of equivalent capacity—both achieve the same σ -RMSE in the ablation study (section 4.5.1). What GatedSigmaNet provides is in-training interpretability: the regime structure of the surface becomes legible during training through the gate function and the two expert components, without any post-hoc clustering, dimensionality reduction, or regime labelling step. The ablation confirms that the full-model gate adjusted Rand index (0.71) exceeds the post-hoc PCA-based clustering of a monolithic surface (0.65), demonstrating that the architectural gate captures regime information incrementally beyond what is recoverable from the surface alone. This property has direct operational value: a risk manager or calibration engineer can inspect the gate curve and the two expert surfaces during training to verify that the model is decomposing the surface in a financially sensible way, without waiting for a full training run to complete and a separate post-hoc analysis to be performed.

The third principal accomplishment is the systematic robustness benchmark presented in section 4.4. The financial calibration literature has lacked a controlled, measurement-consistent head-to-head comparison between a PINN-based inverse calibrator and the Dupire finite-difference method under jointly varying noise and data density. This benchmark’s results illustrate where physics-informed regularization is advantageous. Regularization based on Dupire’s smoothing is comparable with low noise and high data density, and where it is, the PINN advantage is small. As data becomes sparse — the regime of greatest practical interest, covering illiquid single-stock options, pre-open auction books, and stressed market conditions — we observed an RMSE improvement of 4.8 to 38.9 over the full sweep, and Dupire’s R^2 collapsed from 0.997 to 0.342 at 10% data density. These results suggest that, in the cases where calibration quality matters most, physics-informed calibration should be chosen.

5.2. Soft Gating and the Nature of Financial Regime Transitions

The gate ARI of 0.003 for hard binary regime labels requires cautious interpretation. It indicates that the gate learned to assign mixing weights in a continuous manner and not with a binary switch. The switching is not abrupt, and there are no step function transitions in financial volatility regimes. The VIX, which is the primary indicator of equity market volatility, transitions between calm and crisis states over weeks and months. A gate that switches abruptly at a fixed calendar date would impose a discretisation that is not present in the data-generating process; the hyperbolic-tangent transition width of 0.03 in the synthetic data corresponds to a gradual three-to-four-day transition, and a soft gate is the appropriate structural response to this design.

The value of GatedSigmaNet is not in classifying regimes but in decomposing the surface. Expert A and Expert B capture qualitatively distinct smile shapes corresponding to the two volatility regimes without any regime labels, any unsupervised post-processing, or any explicit clustering step. The fact that post-hoc k -means on the recovered synthetic surface

achieves $\text{ARI} = 1.000$ demonstrates that the surface itself encodes the regime information completely. The comparison between the full_pinn gate ARI of 0.71 in the ablation and 0.003 in section 4.2.2 has not been fully reconciled. A likely mechanism is gate over-smoothing: over 20,000 epochs the optimiser drives $g(t)$ toward a monotone sigmoid that minimises the Tikhonov penalty, resulting in low ARI against hard binary labels even when the recovered surface achieves $\text{PCA-ARI} = 1.000$. Both values are reported transparently.

5.3. Practical Implications

Section 4.4 discusses the robustness of the results. Quote densities do not remain consistent across strikes, expiries, or even when market conditions change. In index options, markets are relatively liquid around the ATM; however, the wings may often have no quotes at all. In equity options, quotes may be found for no more than five strikes for the entire duration of the expiry. In times of heightened market stress, when the calibration of volatility surfaces is of utmost importance, bid-ask spreads widen and data density diminishes. Under such circumstances, a calibration method that achieves $R^2 = 0.97$ at 10% data density retains a significant advantage when compared to Dupire's method, which achieves $R^2 = 0.34$ at the same data density.

Practitioners can refer to the Pareto curve (fig. 18) for a principled working value for the physics weight. The Price R^2 profile in the flat region ($\lambda_{\text{PDE}} \in [10^{-2}, 10]$) indicates negligible penalty on price fit, along with monotone improvement in no-arbitrage compliance. Setting $\lambda_{\text{PDE}} = 1$ within this region is both theoretically clean and practically safe. The architecture has rendered the search for the optimal weight irrelevant for volatility and of near-irrelevance for pricing.

5.4. Limitations

Gate interpretability bounds.. The gate ARI in the standalone synthetic experiment is 0.003 against hard binary labels, a considerable drop from the value of 0.71 observed in the ablation at reduced epoch count. This discrepancy has not been fully resolved and constitutes an open problem. Practitioners seeking a hard regime indicator should perform post-hoc k -means on the recovered surface ($\text{ARI} = 1.000$) rather than relying on the gate threshold.

Deep-OTM options on real data.. The SPY test set has an IV RMSE of 0.1135, driven primarily by systematic underfitting at $\text{IV} > 0.40$. This reflects the well-known limitation of local volatility models, which cannot simultaneously fit the near-ATM smile and the deep-OTM wings when those wings are elevated by jump risk premiums outside the scope of a pure diffusion model. Extension to a stochastic-local volatility hybrid [35] or a PINN formulation of the Heston PDE [36] is a natural next step.

At-the-money domain restriction.. All synthetic experiments use at-the-money options ($K = S$) and a restricted spot range $S \in [60, 140]$. Extension to the full moneyness grid is architecturally straightforward but alters the implied-volatility inversion step : far from the money, Brent's method applied to market prices can produce extreme or undefined implied volatilities that contaminate the regression targets. Robust target construction for the full moneyness grid requires additional care.

Computational cost.. A full 20,000-epoch Block-A run requires approximately 2–3 hours on a modern T4 GPU. Physics-informed operator networks [37] and meta-learning approaches [38] are the most promising directions for reducing this cost. For context, the Dupire baseline completes in under 10 seconds on the same hardware; the PINN is approximately 1,000× slower but achieves substantially superior accuracy under sparse and noisy data conditions.

6. Conclusion

We have presented **GatedPINN**, a decoupled inverse Physics-Informed Neural Network framework for local volatility surface calibration. Three design choices jointly distinguish it from prior PINN approaches.

First, the price network and the volatility network are trained with independent optimisers, with the volatility detached from the Black–Scholes PDE residual during price-network training. This detachment severs a fourth-order gradient pathway through the call-price convexity that renders joint training intractable, and has a structural consequence: the volatility RMSE is provably invariant to the PDE weighting coefficient, removing the framework’s most consequential hyperparameter from the calibration workflow.

Second, the local volatility is represented as a continuous mixture of two expert networks blended by a scalar, time-dependent gate. The gate provides in-training architectural decomposition of the volatility surface into two expert components without requiring post-hoc clustering. The experts specialise interpretably to distinct volatility regimes; post-hoc k -means on the recovered synthetic surface recovers the true regime structure with $\text{ARI} = 1.000$.

Third, a systematic head-to-head benchmark against the Dupire finite-difference baseline demonstrates $4.8\times\text{--}38.9\times$ lower RMSE across the noise and sparsity sweeps. The most practically significant result is at 10% data density, where Dupire degrades to $R^2 = 0.34$ and the PINN retains $R^2 = 0.97$, demonstrating that physics-informed regularisation is a viable substitute for data density in sparse market conditions. These results suggest that the combination of decoupled training and architectural regime decomposition is a robust and interpretable approach to the inverse calibration problem. Natural extensions include: (i) replacing the local-volatility Black–Scholes PDE with a stochastic-local volatility or Heston PDE to address deep-OTM fit; (ii) integrating Bayesian PINNs [39] to provide posterior credible intervals on the recovered surface; and (iii) applying physics-informed operator networks [37] to achieve real-time recalibration. We hope the framework and its transparent evaluation, including honest reporting of the gate ARI discrepancy and the deep-OTM limitation, will serve as a useful reference point for these future developments.

Credit authorship contribution statement

Vrinda Bajaj: Conceptualization, Methodology, Software, Formal analysis, Investigation, Writing, original draft, Writing, review & editing. **Veydant Katyal:** Conceptualization, Methodology, Software, Formal analysis, Investigation, Writing, original draft, Writing, review & editing. **Prof. Manoov R:** Supervision, Conceptualization, Writing, review & editing, Resources.

Declaration of Competing Interest

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

Data Availability

The synthetic data generator, training scripts, trained model checkpoints, and analysis notebooks underlying all figures and tables will be made available in a public repository upon acceptance. In the interim, all materials are available from the corresponding author upon reasonable request.

Acknowledgements

The authors are grateful to the School of Computer Science and Engineering at the Vellore Institute of Technology for computational resources and to the open-source communities behind PyTorch, SciPy, and yfinance.

References

- [1] B. Dupire, Pricing with a smile, *Risk* 7 (1) (1994) 18–20.
- [2] E. Derman, I. Kani, Riding on a smile, *Risk* 7 (2) (1994) 32–39.
- [3] T. Hein, B. Hofmann, On the nature of ill-posedness of an inverse problem arising in option pricing, *Inverse Problems* 19 (6) (2003) 1319–1338.
- [4] R. Cont, P. Tankov, *Financial Modelling with Jump Processes*, Chapman & Hall/CRC, London, 2004.
- [5] J. Gatheral, *The Volatility Surface: A Practitioner's Guide*, Wiley Finance, 2006.
- [6] M. Fongler, Arbitrage-free smoothing of the implied volatility surface, *Quantitative Finance* 9 (4) (2009) 417–428.
- [7] J. Andreasen, B. Huge, Volatility interpolation, *Risk* 24 (3) (2011) 86–89.
- [8] R. Lagnado, S. Osher, A technique for calibrating derivative security pricing models: Numerical solution of an inverse problem, *Journal of Computational Finance* 1 (1) (1997) 13–25.
- [9] J. Gatheral, A. Jacquier, Arbitrage-free SVI volatility surfaces, *Quantitative Finance* 14 (1) (2014) 59–71.
- [10] P. Hagan, D. Kumar, A. Lesniewski, D. Woodward, Managing smile risk, *Wilmott Magazine* (2002) 84–108.

- [11] B. Huge, A. Savine, Differential machine learning, *Risk Magazine* (2020).
- [12] J. Cao, J. Chen, J. Hull, Z. Poulos, Deep hedging of derivatives using reinforcement learning, *Journal of Financial Data Science* 3 (1) (2021) 10–27.
- [13] H. Bühler, L. Gonon, J. Teichmann, B. Wood, Deep hedging, *Quantitative Finance* 19 (8) (2019) 1271–1291.
- [14] M. Raissi, P. Perdikaris, G.E. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear PDEs, *Journal of Computational Physics* 378 (2019) 686–707.
- [15] S. Cai, Z. Mao, Z. Wang, M. Yin, G.E. Karniadakis, Physics-informed neural networks (PINNs) for fluid mechanics: A review, *Acta Mechanica Sinica* 37 (2021) 1727–1738.
- [16] E. Haghighat, M. Raissi, A. Moure, H. Gomez, R. Juanes, A physics-informed deep learning framework for inversion and surrogate modeling in solid mechanics, *Computer Methods in Applied Mechanics and Engineering* 379 (2021) 113741.
- [17] K. Linka, A. Schäfer, X. Meng, Z. Zou, G.E. Karniadakis, E. Kuhl, Bayesian physics-informed neural networks for real-world nonlinear dynamical systems, *Computer Methods in Applied Mechanics and Engineering* 402 (2022) 115346.
- [18] A. Salvador, S. Pagliarani, V. Vaccari, Physics-informed neural networks for parametric option pricing, *Proceedings of the AAAI Workshop on Quantitative Finance*, 2020.
- [19] V. Dhiman, J. Hu, Physics-informed neural networks for option pricing, *Quantitative Finance* 23 (10) (2023) 1521–1536.
- [20] P. Bergmann, G. Loeper, S. Ould Aly, Calibration of local volatility models with stochastic interest rates using optimal transport, *Mathematical Finance* 35 (1) (2025) 45–81.
- [21] A. Berardi, A. Cosma, G. Tessitore, Adaptive inverse PINN for local volatility calibration, *Mathematical Finance* 34 (2) (2024) 412–449.
- [22] M. Almanstötter, F. Köhler, J. Pohl, PINNverse: A constrained-optimisation perspective on inverse PINNs, *Proceedings of the International Conference on Machine Learning*, 2025.
- [23] J.-P. Bouchaud, M. Potters, *Theory of Financial Risk and Derivative Pricing*, 2nd ed., Cambridge University Press, 2003.
- [24] L. Bauwens, S. Laurent, J.V.K. Rombouts, Multivariate GARCH models: A survey, *Journal of Applied Econometrics* 21 (2006) 79–109.
- [25] R.A. Jacobs, M.I. Jordan, S.J. Nowlan, G.E. Hinton, Adaptive mixtures of local experts, *Neural Computation* 3 (1991) 79–87.

- [26] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. Le, G. Hinton, J. Dean, Outrageously large neural networks: The sparsely-gated mixture-of-experts layer, *Proceedings of the International Conference on Learning Representations*, 2017.
- [27] D.T. Breeden, R.H. Litzenberger, Prices of state-contingent claims implicit in option prices, *Journal of Business* 51 (1978) 621–651.
- [28] M. Tancik, P. Srinivasan, B. Mildenhall et al., Fourier features let networks learn high-frequency functions in low dimensional domains, *Advances in Neural Information Processing Systems* 33 (2020) 7537–7547.
- [29] N. Rahaman, A. Baratin, D. Arpit et al., On the spectral bias of neural networks, *Proceedings of the International Conference on Machine Learning*, 2019, pp. 5301–5310.
- [30] Y. Bengio, J. Louradour, R. Collobert, J. Weston, Curriculum learning, *Proceedings of the 26th International Conference on Machine Learning*, 2009, pp. 41–48.
- [31] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, *Proceedings of the International Conference on Learning Representations*, 2015.
- [32] L. Hubert, P. Arabie, Comparing partitions, *Journal of Classification* 2 (1985) 193–218.
- [33] S. Osher, M. Burger, D. Goldfarb, J. Xu, W. Yin, An iterative regularization method for total variation-based image restoration, *Multiscale Modeling & Simulation* 4 (2) (2005) 460–489.
- [34] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, Distributed optimization and statistical learning via the alternating direction method of multipliers, *Foundations and Trends in Machine Learning* 3 (2011) 1–122.
- [35] R. Renè, Stochastic-local volatility calibration and simulation using deep learning, *Quantitative Finance* 22 (5) (2022) 881–897.
- [36] S.L. Heston, A closed-form solution for options with stochastic volatility with applications to bond and currency options, *Review of Financial Studies* 6 (1993) 327–343.
- [37] Z. Li, N. Kovachki, K. Azizzadenesheli et al., Fourier neural operator for parametric partial differential equations, *Proceedings of the International Conference on Learning Representations*, 2021.
- [38] C. Finn, P. Abbeel, S. Levine, Model-agnostic meta-learning for fast adaptation of deep networks, *Proceedings of the International Conference on Machine Learning*, 2017.
- [39] L. Yang, X. Meng, G.E. Karniadakis, B-PINNs: Bayesian physics-informed neural networks for forward and inverse PDE problems with noisy data, *Journal of Computational Physics* 425 (2021) 109913.

Appendix A. Hyperparameter Settings

Table A.5 provides the complete hyperparameter settings for all experiments. Settings marked with an asterisk (*) are swept in the corresponding experimental block.

Table A.5: Complete hyperparameter settings. Settings marked * are swept in the corresponding block.

Hyperparameter	Value
PriceNet: hidden layers $\times$ width	5×128 , tanh activation
PriceNet: RFF features / bandwidth	$24 / 0.3$
PriceNet: learning rate (initial $\rightarrow$ final)	$10^{-3} \rightarrow 10^{-5}$ (cosine annealing)
PriceNet: gradient clip norm	1.0
GatedSigmaNet: expert layers $\times$ width	5×128 , tanh activation
GatedSigmaNet: gate layers $\times$ width	2×32 , sigmoid output
GatedSigmaNet: output clamp $[\sigma_{\min}, \sigma_{\max}]$	$[0.05, 0.40]$
GatedSigmaNet: output bias initialisation	$\log(\exp(0.15) - 1) \approx -1.821$
GatedSigmaNet: learning rate (initial $\rightarrow$ final)	$3 \times 10^{-3} \rightarrow 3 \times 10^{-5}$ (cosine annealing)
GatedSigmaNet: gradient clip norm	0.5
λ_{term} (terminal loss weight)	5.0
λ_{PDE} (PDE loss weight)	1.0 (nominal); swept in section 4.5.2*
λ_{smooth} (smoothness weight)	10^{-5}
PDE curriculum warm-up epochs	8,000
Total epochs: Blocks A/B	20,000
Total epochs: Block C (robustness)	10,000
Total epochs: Block D (ablation)	5,000
Batch size: data / PDE collocation	256 / 1,024
Optimiser (both networks)	Adam with cosine annealing

Dear Editor/Reviewer,

Here are the highlights of our manuscript titled “Gated Physics-Informed Neural Networks for Inverse Local Volatility Calibration with Learned Regime Structure”:

- Stop-gradient decoupling removes fourth-order gradient pathology in PINN calibration
- GatedSigmaNet mixture-of-experts resolves volatility into blended functional regimes
- Achieves volatility R^2 of 0.97 at 10% data density vs. Dupire baseline's 0.342
- RMSE reduced 4.8× to 38.9× over Dupire finite-difference baseline across conditions
- Volatility RMSE confirmed λ -invariant across five orders of PDE weighting magnitude

Declaration of Competing Interest

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

Dr. Manoov Rajapandy,
and co-authors Vrinda Bajaj and Veydant Katyal.

Orchid Information

1. Dr. Manoov Rajapandy: <https://orcid.org/0000-0003-4745-3122>
2. Vrinda Bajaj: <https://orcid.org/0009-0007-5829-2703>
3. Veydant Katyal: <https://orcid.org/0009-0003-1878-5221>