OXFORD JOURNAL
OF STUDENT
SCHOLARSHIP

Fatigue-Modulated Heston Pricing with Physics-Informed Neural Networks

Ethan Li

ethanli55555@gmail.com

ABSTRACT

Reliable option-pricing research must separate improvements in numerical learning from evidence that a proposed mechanism describes markets. This study tests whether a cumulative, fatigue-inspired shock index can usefully modify the variance target in a synthetic Heston model. European calls on a non-dividend-paying spot asset are priced with neural networks trained using both simulated prices and the governing pricing equation. An independent validation sample selects the equation-loss weight before two fresh test samples are generated. A conventional Fourier solution for the same time-dependent model supplies the evaluation reference and is checked against separate Monte Carlo calculations. Across eight training seeds, the selected network reduces pricing error relative to the data-only network by 10.3% on a fixed-variance test and 12.7% on a variance-spanning test. Enforcing the generating fatigue strength gives a small advantage over plain Heston physics on the variance-spanning sample, but the fixed-variance sample favors plain Heston physics. Conventional numerical pricing remains more accurate than the neural approximations. The findings demonstrate conditional sensitivity to a prescribed synthetic correction; they do not establish a fatigue mechanism in real markets or predictive validity for energy futures.

1 INTRODUCTION

Reliable option-pricing experiments must distinguish a useful model assumption from an accurate implementation of that assumption. A neural network may benefit from a pricing equation because the equation supplies information missing from a small training sample. That benefit alone does not show that an additional mechanism describes financial markets. This distinction matters when a familiar mathematical model is extended using an analogy from another field.

The Heston model represents volatility through a random variance process that returns toward a long-run level [1]. A physics-informed neural network (PINN) approximates the option price while penalizing violations of the pricing equation [2]. Here, “physics” means the stipulated mathematical dynamics; it does not imply a physical law governing asset prices. Hainaut and Casas applied this approach to conventional Heston pricing [3]. Our question concerns the effect of modifying the variance target and then enforcing the resulting equation during training.

We consider a cumulative index inspired by the power-law relationship between loading amplitude and fatigue life in the Coffin–Manson tradition [4,5], together with the cumulative-damage idea associated with Miner [6]. Larger synthetic shock marks produce larger increments in this index. The index raises the variance target over time. This is a chosen functional analogy, not a derivation of financial dynamics from material mechanics. Once fixed, the index is a deterministic schedule, so the resulting model remains a Heston model with a time-dependent coefficient. Such models also admit conventional transform-based numerical pricing [7].

The study asks three questions. Does enforcing the specified pricing equation improve a network trained on noisy simulated prices? How does that network compare with conventional numerical pricing under the same model? Does enforcement of the generating fatigue strength outperform enforcement of two alternative strengths on fresh contracts? The third question is a restricted sensitivity experiment, not a continuous estimation of an unknown parameter.

The revised experiment uses a non-dividend-paying **spot asset**. The original submission called the underlying an energy futures price even though its risk-neutral drift was the spot drift. That interpretation is corrected throughout. No energy option-chain data, convenience yields, storage costs, delivery mechanics, or futures settlement conventions are modeled. All conclusions concern the specified synthetic spot-option experiment.

2 METHODS

2.1 MODEL AND MEANING OF THE CORRECTION

Let S denote the spot asset price, v its instantaneous variance, and t calendar time measured in model years. All options are European calls with strike $K = 3.20$ and common expiry $T = 1$. The short rate is constant at $r = 0.045$. Under the risk-neutral measure, the asset drift is rS and variance follows the Heston square-root process. Appendix A gives the stochastic equations and the complete pricing equation.

The extension replaces the constant long-run variance with

$$\theta_{eff}(t) = \theta_0[1 + \beta D(t)]. \qquad (1)$$

Here $\theta_0 = 0.09$, $D(t)$ is the specified damage index, and β controls its influence. The generating value is $\beta = 0.5$, so the variance target increases from 0.09 to 0.135 over the schedule. Setting $\beta = 0$ gives conventional Heston dynamics. The mean-reversion speed remains $\kappa = 2$; the correction changes its target, not the speed itself. Volatility of variance is $\xi = 0.35$ and the correlation of the Brownian shocks is $\rho = -0.55$.

The deterministic schedule adds no stochastic state. The option price is a function $C(S, v, t)$, and the pricing equation is linear in C. Prices are continuous across changes in the prescribed variance target, although their time derivative can change. An endogenous damage process based on future realized asset movements would require a separate model and potentially additional state or jump terms; it is not implemented here.

2.2 DAMAGE CONSTRUCTION AND THE CALENDAR MAPPING

The released generator sets its step length to $T/1260$. Consequently its 1,260 steps span **one model year**, despite the earlier description of approximately five trading years. The revision follows the implemented clock. On each step, an independent Bernoulli event occurs with probability 0.01. Conditional on an event, a log-jump mark is drawn from a normal distribution with mean zero and standard deviation 0.12. The seed-0 realization contains 13 nonzero marks. These are synthetic event marks, not detected excursions of the complete simulated return series, and the event count is binomial rather than exactly Poisson.

To preserve the original seed-0 marks, the generator consumes its two diffusion-normal draws before the event draw on each step. Only the marks are used to define the schedule; the simulated asset trajectory used to create them is not reused as an option-pricing path. For grid index $j = 0, \ldots, 1260$, define

$$a_j = 1_{\{|J_j|>\delta\}}|J_j|^{-1/c}, \qquad D_j = \frac{\sum_{k=0}^{j} a_k}{\sum_{k=0}^{1260} a_k}, \qquad c = -0.6, \quad \delta = 10^{-6}. \qquad (2)$$

If the denominator is zero, all D_j are set to zero. Because $-1/c > 0$, a larger mark contributes more damage. The common fatigue amplitude scale cancels when cumulative damage is divided by its terminal value; it is therefore not an independently identifiable parameter in this normalized experiment. The small denominator stabilizer in the original implementation is replaced by the explicit zero-damage case.

The calendar rule is

$$D(t) = D_{min(\lfloor 1260t \rfloor, 1260)}, \qquad 0 \le t \le 1. \qquad (3)$$

Every contract uses the same schedule. A contract observed at $t = 0.40$ starts with D_{504} and uses the remaining schedule from 0.40 to 1.00; its remaining maturity is 0.60 years. The schedule is not restarted or stretched for that contract. In the training Monte Carlo calculation, each step evaluates the schedule at the step's starting calendar time. The Fourier benchmark integrates each constant segment exactly before numerical inversion.

The entire schedule, including its terminal normalization, is assumed known in this synthetic world. This is a conditional pricing experiment, not a procedure that forecasts future shocks from past data. A real-time application could not normalize by unobserved future shocks. Figure 1 shows the prescribed marks, cumulative damage, and variance target.

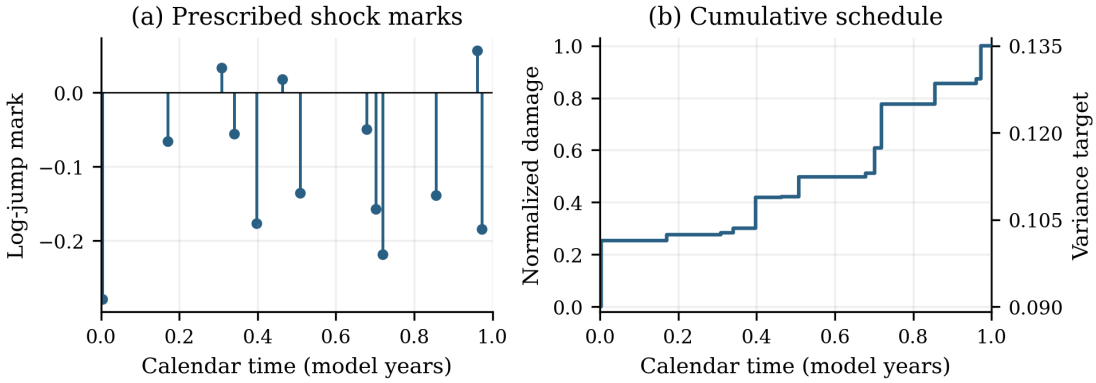


Figure 1. Prescribed one-year damage schedule. (a) The 13 seed-0 log-jump marks. (b) The cumulative normalized index and the corresponding variance target at generating $\beta = 0.5$. Every contract uses its remaining segment of this same calendar schedule.

2.3 CONVENTIONAL NUMERICAL REFERENCE AND BASELINES

We implement a characteristic-function pricer for the same deterministic, piecewise-constant variance target. The affine Heston representation permits analytic integration of the variance-target contribution over each segment, followed by one-dimensional numerical Fourier inversion [1,7]. Appendix A provides the formulas. This supplies a conventional numerical benchmark for the fatigue-modulated model itself, not merely for the fatigue-free model.

The base inversion uses Gauss–Legendre quadrature with a contract-specific cutoff and at least 512 nodes. A refined inversion increases the cutoff by 50% and doubles the nodes. Refined prices are used for validation and final evaluation, without added observation noise. The base-versus-refined difference measures numerical refinement sensitivity; it is not an independently certified error bound. Both share the same characteristic function, so additional checks use the deterministic-variance limit, the pricing-equation residual evaluated by finite differences, and an independently coded Monte Carlo simulator.

The Monte Carlo check uses 120,000 paths at each of six contract settings, with nominal grids of 200 and 800 steps. All damage-change times are inserted into those grids. Asset updates use beginning-of-step variance, and the variance update uses full truncation [8]. Monte Carlo standard

errors describe payoff sampling variability, not discretization error. Testing two grids helps assess the latter, but does not prove its absence.

Comparators are analytic Black–Scholes with constant volatility 0.30 [9], conventional Heston Fourier pricing with $\beta = 0$, the matched fatigue-modulated Fourier benchmark, and the neural models below. All use the same rate, strike, valuation time, and spot convention. The plain Heston numerical baseline receives each contract's actual initial variance. These are controlled parameter settings, not independently calibrated market models. Including the matched numerical solution makes the distinction between approximation error and model mismatch explicit.

2.4 DATA SEPARATION AND NETWORK TRAINING

Table 1 summarizes the data protocol. The training sample has 600 contracts and uses the original corrected Monte Carlo labeling approach: 4,000 paths, 50 steps, and independent observation noise with standard deviation 0.005. Each simulated path starts from its contract's sampled variance. This deliberately retains noisy training supervision while evaluating against a substantially more accurate reference. Finite-step bias and sampling noise remain possible in the training labels.

Table 1. Model and data settings for the corrective experiment. Uniform inputs are sampled independently within each dataset. All expiries are T = 1; remaining maturity is 1 − t.

Setting	Specification
Spot model	$K = 3.20$; $r = 0.045$; $\kappa = 2$; $\theta_0 = 0.09$; $\xi = 0.35$; $\rho = -0.55$
Damage schedule	1,260 steps over one year; seed 0; 13 events; generating $\beta = 0.5$
Shared input ranges	$S \sim U(2.24, 4.48)$; $t \sim U(0.05, 1)$; varying $v \sim U(0.03, 0.25)$
Training	600 contracts; seed 1; 4,000 Monte Carlo paths and 50 steps; noise SD 0.005
Validation	400 variance-spanning contracts; seed 2026092401; refined Fourier prices
Slice test	400 contracts; $v = 0.09$; seed 2026092402; refined Fourier prices
Span test	400 variance-spanning contracts; seed 2026092403; refined Fourier prices
Training randomness	Eight optimizer seeds, 42–49, for each configuration

The independent validation set contains 400 contracts spanning the training variance range. For each candidate physics weight $w_p \in \{0.01, 0.1, 1\}$, eight models are trained using optimizer seeds 42–49 and the generating $\beta = 0.5$. The selected weight minimizes mean validation root mean squared error (RMSE) across those seeds; a tie would favor the smaller weight. The selection is saved before either new test set is generated. The originally reported slice test is not used for this decision.

Two fresh test samples contain 400 contracts each. The slice sample fixes $v = 0.09$, while the span sample varies v over the training range. All other input ranges are shared. The tests assess interpolation in different variance geometries, not extrapolation. The sample sizes, random seeds, candidate weights, parameter arms, selection rule, and statistical comparisons were recorded before the corrective rerun. This is a revision protocol, not a claim that the original study was preregistered.

The network has three inputs, four hidden layers of 64 hyperbolic-tangent units, and one price output. Adam training runs for 1,500 epochs with initial learning rate 10^{-3} and cosine decay. The loss is

$$L = L_{data} + w_p L_{PDE} + L_{terminal}. \qquad (4)$$

The first term is mean squared error against training labels. The second is the mean squared pricing-equation residual computed by automatic differentiation at 2,000 newly sampled points per epoch. The third enforces the call payoff at expiry using 500 newly sampled points. Collocation and terminal samples cover $S \in [0.96, 8.00]$ and $v \in [0.001, 0.5]$; collocation time covers $[0, 1]$. No separate spatial-boundary loss, no-arbitrage penalty, or hard output constraint is imposed. These choices limit the claim to the evaluated domain.

The data-only ablation sets $w_p = 0$ while retaining terminal supervision. After selection, two additional groups enforce $\beta = 0$ or $\beta = 2$ at the selected weight. There are 48 trained models: eight data-only, 24 in the weight sweep, and 16 additional parameter-arm models. The matched arm reuses its sweep models. Each neural configuration uses the same eight initialization seeds and the same fixed training sample. Matched and alternative physics arms also share the sampling sequence under a given seed.

2.5 OUTCOMES AND STATISTICAL INTERPRETATION

The primary outcome is RMSE against refined numerical prices on all 400 contracts. The secondary outcome is mean absolute percentage error (MAPE), restricted to reference prices above 0.05 and $|S/K - 1| < 0.4$ to reduce instability near zero prices. Tables report the mean and **sample standard deviation** across eight training seeds; full ranges and per-seed records are supplied electronically. Numerical baselines are deterministic and consequently have no training-seed standard deviation.

For each test geometry, we compare the matched PINN's RMSE with the data-only, plain-physics, and overshot-physics networks. Differences are paired by training seed. We report the mean difference, an

unadjusted 95% Student t interval with seven degrees of freedom, the number of improving seeds, and a two-sided p value adjusted by Holm's procedure across these six comparisons [10]. The intervals are not simultaneous confidence intervals. MAPE is reported descriptively rather than introducing a second family of significance claims.

The eight seeds quantify variation from training the network conditional on one shock schedule and one shared set of training and evaluation contracts. They do not represent eight independent markets or label-generation experiments. With only eight observations per comparison, inference also depends on the suitability of the paired-difference distribution for the t procedure. Effect sizes and their practical scale therefore accompany the significance calculations. Failure to reject a difference is not evidence of equivalence.

The submitted trading proxy used adjacent, randomly ordered contracts as if they formed a time sequence. It has been removed from the reported outcomes because it cannot establish realizable trading performance. The revision evaluates pricing accuracy only.

3 RESULTS

3.1 VERIFICATION OF THE NUMERICAL REFERENCE

Across the 800 test contracts, increasing the Fourier cutoff and node count changes prices by at most 1.93×10^{-12}. The base-versus-refined RMSE is 1.27×10^{-13} on the slice and 8.15×10^{-14} on the span. The constant-variance limit agrees with analytic Black–Scholes to 3.60×10^{-13} over four check contracts. Finite-difference residuals of the pricing equation at three interior points, away from schedule changes, have maximum absolute value 2.75×10^{-8}.

All 12 Monte Carlo check prices lie within 1.25 estimated sampling standard errors of the Fourier reference (Figure 2). The discounted-asset checks deviate from their martingale expectation by at most 2.48 sampling standard errors across these runs. On the test contracts, reference prices satisfy the standard spot-call bounds to within 2×10^{-12} numerical tolerance. Taken together, these are consistency checks on the implementation; quadrature refinement alone would not provide an independent validation.

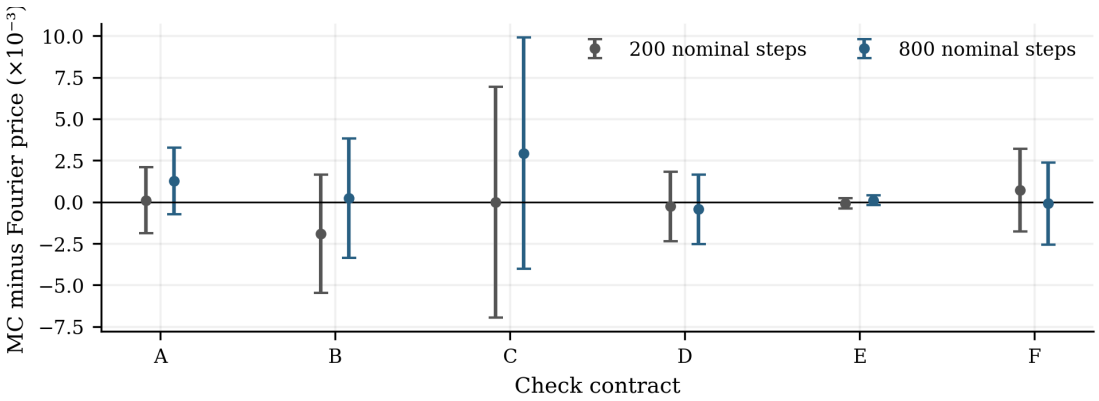


Figure 2. Independent simulation checks against the matched Fourier reference. Each point uses 120,000 paths; bars are ±1.96 Monte Carlo standard errors and exclude discretization uncertainty. Both nominal grids also include every damage-change time. Contract settings (S, v, t) are A = (2.8, 0.04, 0.10), B = (3.2, 0.09, 0.10), C = (4.0, 0.20, 0.10), D = (3.2, 0.09, 0.65), E = (2.8, 0.04, 0.80), and F = (3.2, 0.25, 0.80).

3.2 VALIDATION SELECTION

Mean validation RMSE is 0.03522 at weight 0.01, 0.03176 at 0.1, and 0.03122 at 1.0 (Table 2). The prespecified minimum-mean rule selects $w_p = 1.0$. This is an operational hyperparameter choice, not a claim that the difference between the two leading validation means is statistically significant. The selected value is then held fixed for both test geometries and both alternative coupling strengths.

Table 2. Validation RMSE (×10⁻³) across eight training seeds. Only positive weights are candidates for the matched-physics selection.

Physics weight	Mean ± sample SD	Observed range	Selected
0	35.077 ± 2.377	32.284–38.590	Ablation only
0.01	35.216 ± 2.189	31.924–39.044	No
0.1	31.764 ± 2.694	27.850–36.045	No
1	31.219 ± 1.026	29.576–32.925	Yes

3.3 FRESH TEST PERFORMANCE

Table 3. Fresh test performance against refined Fourier prices. RMSE is shown in units of 10⁻³; MAPE is a percentage. Neural entries are mean ± sample SD over eight seeds. All PINNs use selected weight 1.0.

Model	Slice RMSE	Slice MAPE	Span RMSE	Span MAPE
Black–Scholes	14.203	2.62	45.443	12.60
Plain Heston Fourier	10.939	3.92	9.614	3.40
Matched Fourier*	1.27×10^{-10}	< 0.000001	8.15×10^{-11}	< 0.000001
Data-only network	34.280 ± 1.931	9.53 ± 1.11	36.419 ± 2.196	9.97 ± 1.14
PINN, β = 0	30.495 ± 1.121	5.56 ± 0.43	32.146 ± 1.140	7.28 ± 0.52
PINN, β = 0.5	30.761 ± 1.048	5.75 ± 0.36	31.805 ± 1.103	6.97 ± 0.50
PINN, β = 2	34.436 ± 1.025	7.98 ± 0.28	33.896 ± 1.072	7.60 ± 0.27

*Matched Fourier entries measure the base-versus-refined inversion difference within the same solver. They are refinement diagnostics, not a certified absolute pricing error. Deterministic numerical methods have no seed SD.

The matched PINN reduces slice RMSE from 0.03428 ± 0.00193 to 0.03076 ± 0.00105 and span RMSE from 0.03642 ± 0.00220 to 0.03181 ± 0.00110 relative to the data-only network (Table 3). These are relative reductions of 10.3% and 12.7%. RMSE is lower in all eight paired seeds on both tests. Holm-adjusted p values are 0.00475 and 0.00193, respectively (Table 4). MAPE also decreases, from 9.53% to 5.75% on the slice and from 9.97% to 6.97% on the span. The MAPE filters retain 323 and 332 contracts.

Black–Scholes is more accurate than the matched PINN on the fixed-variance slice, but its span RMSE is 0.04544. The matched PINN has about 30.0% lower span RMSE than this constant-volatility comparator. Conventional plain Heston Fourier pricing, however, has still lower RMSE: 0.01094 on the slice and 0.00961 on the span. The matched Fourier solver agrees with its refined version to the numerical scale reported above. Thus, the new comparison supports a benefit from equation regularization within the neural architecture, while showing that conventional numerical pricing remains substantially more accurate for these contracts.

3.4 SENSITIVITY TO THE ENFORCED FATIGUE STRENGTH

On the span sample, matched physics has lower mean RMSE than plain Heston physics by 0.000341 (95% paired interval −0.000505 to −0.000177; Holm-adjusted $p = 0.00475$). This is a 1.06% reduction, with improvement in seven of eight seeds. Mean span MAPE is also lower: 6.97% versus 7.28%. Both measures favor the generating strength over the overshot strength as well. These comparisons support limited sensitivity to the prescribed correction in the variance-spanning setting.

The slice sample gives a different result. Matched physics has **higher** mean RMSE than plain physics by 0.000266 (95% interval 0.000057 to 0.000475; Holm-adjusted $p = 0.01963$), and mean MAPE is 5.75% versus 5.56%. Plain physics has lower RMSE in seven of eight seeds. The revised slice result therefore favors the simpler physics; it is not an equivalence or an exact tie. Enforcing $\beta = 2$ gives higher mean RMSE than matched physics on both tests. Figure 3 displays this dependence together with the independent validation sweep.

Table 4. Paired RMSE contrasts: matched PINN minus the indicated neural comparator. Differences and confidence limits use units of 10^{-3}. Negative differences favor matched physics. Intervals are unadjusted 95% t intervals; p values use Holm adjustment across all six rows.

Test	Comparator	Mean difference	95% interval	Holm p	Improved seeds
Slice	Data-only	-3.520	[-5.187, -1.852]	0.00475	8/8
Slice	β = 0	+0.266	[+0.057, +0.475]	0.01963	1/8
Slice	β = 2	-3.675	[-4.223, -3.127]	$5.76 \times 10^{-}$	8/8
Span	Data-only	-4.614	[-6.398, -2.831]	0.00193	8/8
Span	β = 0	-0.341	[-0.505, -0.177]	0.00475	7/8
Span	β = 2	-2.090	[-2.562, -1.619]	$7.85 \times 10^{-}$	8/8

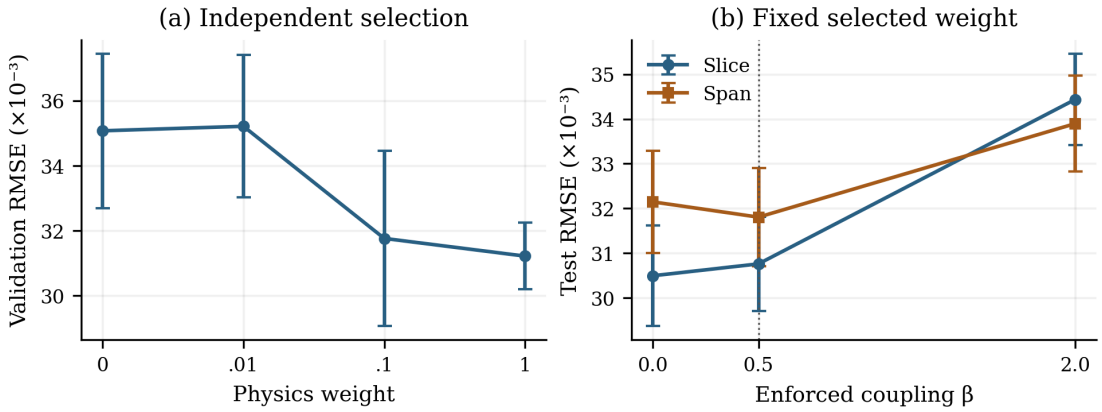


Figure 3. (a) Validation RMSE versus physics weight; weight 0 is the data-only ablation and is not a candidate for selecting the matched PINN. (b) Fresh-test RMSE versus enforced coupling at the selected weight 1.0. Points show eight-seed means and bars show sample SD, not confidence intervals. The dotted line identifies the generating β = 0.5. Coupling values are unevenly spaced.

4 DISCUSSION

4.1 WHAT EQUATION REGULARIZATION CONTRIBUTES

The matched equation improves the network relative to the data-only ablation on both new test sets. The effect is larger than the extra benefit associated with the fatigue term on the span sample. That distinction is central: a 12.7% improvement from equation regularization relative to data-only training coexists with only a 1.06% improvement from matched versus plain Heston physics on that sample. On the slice, the fatigue term modestly worsens neural RMSE. The evidence supports conditional learning benefits and a small, geometry-dependent sensitivity to the correction.

The conventional benchmark changes the interpretation of neural performance. The matched Fourier calculation directly solves the specified synthetic model and has negligible refinement sensitivity at the tested points. Neural residual error is therefore an approximation cost in this experiment. A PINN may be useful as a learned surrogate or in settings with additional data constraints, but this study does not demonstrate a speed advantage, a cost advantage, or superiority over conventional pricing. Training and inference times were not compared under a controlled performance protocol.

Black–Scholes is a useful simple comparator, but its constant volatility cannot use the varying initial-variance input. Its changing performance across the slice and span samples is consequently conditional on the evaluation distribution. Neither a favorable comparison with Black–Scholes nor an advantage over the data-only network isolates the contribution of fatigue modulation. The paired $\beta = 0$ versus $\beta = 0.5$ neural comparison addresses that narrower question, and the conventional plain Heston baseline shows the scale of price changes due to the prescribed schedule.

4.2 LIMITS OF THE FATIGUE INTERPRETATION

The damage analogy supplies a monotone, power-law transformation of synthetic marks. It does not identify a physical source of market volatility. Because future damage is specified deterministically, this experiment examines a time-dependent Heston coefficient rather than realized-path feedback. Only one schedule, exponent, generating strength, parameter configuration, and strike are studied. Normalization removes the amplitude-scale parameter, so recovering that scale jointly with the coupling strength is not supported by this formulation.

The three enforced strengths are also unevenly spaced: their distances from the generating value are 0.5 and 1.5. A larger penalty at $\beta = 2$ would therefore not establish asymmetric sensitivity. Comparing three predefined values does not identify a unique optimum over continuous β, and network approximation error can obscure or reverse small differences between their pricing implications. The revised results are reported as conditional sensitivity evidence rather than a general recovery theorem.

4.3 LIMITATIONS AND NEXT EXPERIMENTS

The training labels retain Monte Carlo noise, observation noise, and potential finite-step bias. The numerical reference was checked through refinement, limiting cases, the differential equation, and separate simulation, but these checks do not establish a rigorous global error bound. Eight training seeds and two fresh test samples leave uncertainty about other data draws and schedules. All tests are within the sampled input ranges; no inference is made about extreme moneyness, variance extrapolation, other expiries or strikes, or arbitrage consistency outside the tested points.

Further work should first repeat the experiment across independently generated damage schedules and training-label samples, separating those sources of variation from optimizer randomness. A denser, balanced coupling-strength grid and an actual inverse-estimation objective would be needed to investigate identification. An endogenous damage model would require a causal construction and a consistent expanded state description. Any application to energy futures would additionally require appropriate futures dynamics, contract conventions, real option-chain calibration, and comparisons with conventional calibrated models. None of those validations is supplied by the present synthetic results.

5 CONCLUSION

A deterministic fatigue-inspired schedule can be incorporated consistently into Heston spot-option pricing and used as a neural training constraint. With independent validation and fresh tests, matched equation regularization improves RMSE over the data-only network on both evaluation geometries. The extra contribution of matched fatigue physics is small and conditional: it helps on the variance-spanning sample, while plain physics performs better on the fixed-variance slice. Conventional numerical Heston pricing remains more accurate than the neural approximations. These results establish a reproducible synthetic sensitivity experiment. They do not validate a fatigue mechanism in real financial markets, identify a continuous coupling parameter, or demonstrate predictive performance for energy futures.

DATA AND CODE AVAILABILITY

Original source code is available at https://github.com/apersonli55555/fatigue-pinn-option-pricing. The accompanying revision archive contains the corrective experiment code and data, frozen protocol, 48 trained models, statistical summaries, and figure sources. It includes instructions to audit saved predictions or repeat the full run. Exact floating-point agreement across hardware or software versions is not guaranteed. The archive supplies the revised results; these revision files have not been added to the public repository.

REFERENCES

[1] Heston, S. L. "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options." *The Review of Financial Studies*, 6(2), 327–343, 1993. https://doi.org/10.1093/rfs/6.2.327.

[2] Raissi, M., 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. https://doi.org/10.1016/j.jcp.2018.10.045.

[3] Hainaut, D., and A. Casas. "Option Pricing in the Heston Model with Physics Inspired Neural Networks." *Annals of Finance*, 20, 353–376, 2024. https://doi.org/10.1007/s10436-024-00452-7.

[4] Coffin, L. F., Jr. "A Study of the Effects of Cyclic Thermal Stresses on a Ductile Metal." *Transactions of the ASME*, 76(6), 931–949, 1954. https://doi.org/10.1115/1.4015020.

[5] Manson, S. S. *Thermal Stress and Low-Cycle Fatigue*. McGraw-Hill, 1966.

[6] Miner, M. A. "Cumulative Damage in Fatigue." *Journal of Applied Mechanics*, 12(3), A159–A164, 1945.

[7] Elices, A. "Models with Time-Dependent Parameters Using Transform Methods: Application to Heston's Model." arXiv:0708.2020v2, revised 2008. https://arxiv.org/abs/0708.2020.

[8] Lord, R., R. Koekkoek, and D. van Dijk. "A Comparison of Biased Simulation Schemes for Stochastic Volatility Models." *Quantitative Finance*, 10(2), 177–194, 2010. https://doi.org/10.1080/14697680802392496.

[9] Black, F., and M. Scholes. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, 81(3), 637–654, 1973. https://doi.org/10.1086/260062.

[10] Holm, S. "A Simple Sequentially Rejective Multiple Test Procedure." *Scandinavian Journal of Statistics*, 6(2), 65–70, 1979. https://www.jstor.org/stable/4615733.

APPENDIX A STANDARD DYNAMICS AND THE NUMERICAL SOLUTION

A1 RISK NEUTRAL SPOT DYNAMICS

The stochastic differential equations are

$$dS_t = rS_t\,dt + \sqrt{v_t}S_t\,dW_t^{(1)}, \qquad (A1)$$

$$dv_t = \kappa\left[\theta_{eff}(t) - v_t\right]dt + \xi\sqrt{v_t}\,dW_t^{(2)}, \qquad dW_t^{(1)}dW_t^{(2)} = \rho\,dt. \qquad (A2)$$

The two Brownian motions represent correlated price and variance innovations. All coefficients are risk-neutral model inputs. The deterministic schedule fixes the variance target at each calendar time. This is a spot model with no payout; interpreting the same variable as a futures price would not justify the rS drift under the stated constant-rate setting.

For a call with payoff $\left(S_T - K\right)^+$, the pricing equation is

$$C_t + \tfrac{1}{2}vS^2C_{SS} + \rho\xi vSC_{Sv} + \tfrac{1}{2}\xi^2vC_{vv} + rSC_S + \kappa\left[\theta_{eff}(t) - v\right]C_v - rC = 0. \qquad (A3)$$

Here subscripts denote partial derivatives. The second-derivative terms describe diffusion in price and variance and their correlation. The first-derivative terms describe their risk-neutral drift, while $-rC$ accounts for discounting. The terminal condition is $C(S, v, T) = (S - K)^+$. Applying Itô's formula to the discounted price and setting its drift to zero gives (A3). The PDE residual in the neural loss is the left-hand side of this equation.

A2 PIECEWISE VARIANCE TARGET AND FOURIER INVERSION

Let $\tau = T - t$ and define the characteristic function of the log return by $\phi(u) = E\left[exp\left(iulog\left(S_T/S_t\right)\right)|S_t, v_t\right]$. For the stated constant κ, ξ, ρ, write

$$\phi(u) = exp\left(iur\tau + \kappa\int_t^T \theta_{eff}(s)B(u, T - s)\,ds + B(u,\tau)v\right). \qquad (A4)$$

The standard affine calculation yields

$$b = \kappa - \rho\xi iu, \qquad d = \sqrt{b^2 + \xi^2\left(u^2 + iu\right)}, \qquad g = \frac{b-d}{b+d}, \qquad (A5)$$

$$B(u, a) = \frac{b-d}{\xi^2}\frac{1-e^{-da}}{1-ge^{-da}}, \qquad (A6)$$

$$J(u,a) = \int_0^a B(u,z)\,dz = \frac{(b-d)a - 2log\left[\left(1-ge^{-da}\right)/(1-g)\right]}{\xi^2}. \qquad (A7)$$

The square root is chosen with nonnegative real part. The logarithm is evaluated in the stable decaying-exponential representation above for this parameter domain. An interval $[a, b]$ on which the target is constant at θ_j contributes $\theta_j[J(u, T - max(t, a)) - J(u, T - b)]$ to the integral in (A4), provided $b > max(t, a)$. Intervals before valuation are omitted. This integrates the time dependence without stepping a pricing PDE or simulating paths.

Let $k = log(S/K)$. The call price is $SP_1 - Ke^{-r\tau}P_2$, where

$$P_2 = \frac{1}{2} + \frac{1}{\pi}\int_0^\infty Re\left[\frac{e^{iuk}\phi(u)}{iu}\right]du, \qquad (A8)$$

$$P_1 = \frac{1}{2} + \frac{1}{\pi}\int_0^\infty Re\left[\frac{e^{iuk}\phi(u-i)}{iuexp(r\tau)}\right]du. \qquad (A9)$$

In the base computation the integration cutoff is $U = max\left(200, 18/\sqrt{v\tau}\right)$. The node count is $n = max(512, 128\lceil 0.7U|k|/128\rceil + 256)$. The reference calculation uses $1.5U$ and $2n$. At expiry the code returns the payoff directly. Source code records the deterministic-variance limit separately to avoid division by zero at $\xi = 0$.

APPENDIX B RELATION TO THE SUBMITTED EXPERIMENT

The earlier development found an update-order error in a Monte Carlo generator and a separate failure to initialize paths at each contract's supplied variance. The corrected generator uses beginning-of-step variance for the asset update and the contract-specific initial variance. Those corrections are retained. The extended narrative of the discovery and the obsolete plots are omitted from the revised main text so that verification of the current experiment takes priority.

The revision also corrects the spot/futures terminology and one-year clock, replaces test-based weight selection with independent validation, creates fresh test inputs, and adds a numerical reference for the actual time-dependent model. Its reported numerical results therefore replace, rather than extend, the submitted tables. The original 41% comparison with Black–Scholes is not carried forward as an estimate of the fatigue contribution. The unrelated-contract trading statistic is removed. Complete seed-level values remain in the accompanying data files, while the manuscript presents aggregate stability measures.