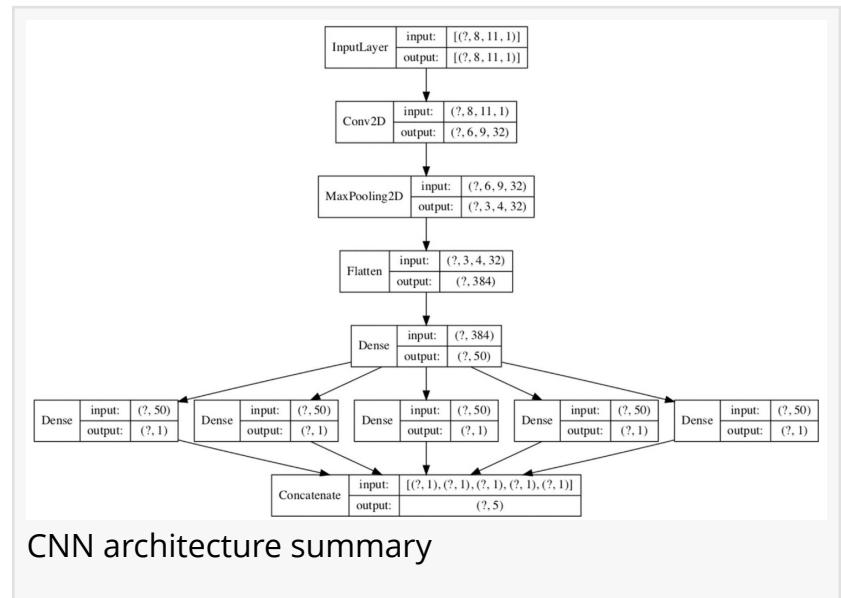


# Interpretability in deep learning for finance: A case study for the Heston model

GA, UNITED KINGDOM, January 29, 2026 /EINPresswire.com/ -- [Deep learning](#) is increasingly used in financial modeling, but its lack of transparency raises risks. Using the well-known Heston option pricing model as a benchmark, researchers show that global interpretability tools based on Shapley values can meaningfully explain how neural networks work, outperforming local explanation methods and helping identify better-performing network architectures.



CNN architecture summary

Deep learning has become a powerful tool in quantitative finance, with applications ranging from option pricing to model calibration. However, despite its accuracy and speed, one major concern remains: neural networks often behave like “black boxes”, making it difficult to understand how they reach their conclusions. This spells a lack of validation, accountability, and risk management in financial decision-making.

In a new study published in *Risk Sciences*, a team of researchers from Italy and the UK investigate how interpretable deep learning models can be made in a financial setting. Their goal was to understand whether interpretability tools can genuinely explain what a neural network has learned, rather than just producing visually appealing but potentially misleading explanations.

The researchers focused on the calibration of the Heston model, one of the most widely used stochastic volatility models in option pricing, whose mathematical and financial properties are well understood. This makes it an ideal benchmark for testing whether interpretability methods provide explanations that align with established financial intuition.

“We trained neural networks to learn the relationship between volatility smiles and the underlying parameters of the Heston model, using synthetic data generated from the model

itself," shares lead author Damiano Brigo, a professor of mathematical finance at Imperial College London. "We then applied a range of interpretability techniques to explain how the networks mapped inputs to outputs."

These techniques included local methods—such as LIME, DeepLIFT, and Layer-wise Relevance Propagation—as well as global methods based on Shapley values, originally developed in cooperative game theory.

The results showed a clear distinction between local and global interpretability approaches. "Local methods, which explain individual predictions by approximating the model locally, often produced unstable or financially unintuitive explanations," says Brigo. "In contrast, global methods based on Shapley values consistently highlighted input features—such as option maturities and strikes—in ways that aligned with the known behavior of the Heston model."

The team's analysis also revealed that Shapley values can be used as a practical diagnostic tool for model design. By comparing different neural network architectures, the researchers found that fully connected neural networks outperformed convolutional neural networks for this calibration task, both in accuracy and interpretability—contrary to what is commonly observed in image recognition.

"Shapley values not only help explain model predictions, but also help us choose better neural network architectures that reflect the true financial structure of the problem," explains co-author Xiaoshan Huang, a quantitative analyst at Barclays.

By demonstrating that global interpretability methods can meaningfully reduce the black-box nature of deep learning in finance, the study provides a pathway toward more transparent, trustworthy, and robust machine-learning tools for financial modeling.

DOI

[10.1016/j.risk.2025.100030](https://doi.org/10.1016/j.risk.2025.100030)

Original Source URL

<https://doi.org/10.1016/j.risk.2025.100030>

Lucy Wang

BioDesign Research

[email us here](#)

---

This press release can be viewed online at: <https://www.einpresswire.com/article/887606639>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable

in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.