Explainable AI Methods in Finance

Quant Strats EU 2025

Daniele Bianchi

Queen Mary, University of London
School of Economics and Finance

Today's Talk

1. **The Business Case:** Why explainability matters
2. **Explainable AI (XAI) Tools:** SHAP and Surrogate Models you can use from today
3. **Actionable Insights:** How to implement and what to expect at deployment

Why Explainability Matters

Why Your Best Models Aren't Deployed

The reality:

- Your quants build models that outperform in backtests
- Risk committees won't approve them
- Compliance blocks deployment

Why Your Best Models Aren't Deployed

The reality:

- Your quants build models that outperform in backtests
- Risk committees won't approve them
- Compliance blocks deployment

The problem: "How do you explain a £100M position when your model can't?"

Why Your Best Models Aren't Deployed

The reality:

- Your quants build models that outperform in backtests
- Risk committees won't approve them
- Compliance blocks deployment

The problem: "How do you explain a £100M position when your model can't?"

The actual cost of model opacity:

- Development costs wasted
- Slower time-to-market vs. firms with explainable models
- Lost AUM to more transparent competitors

The Regulatory Walls Are Closing In

The regulatory landscape is changing:

- **EU AI Act (2024):** High-risk AI requires transparency and explainability
- **MiFID II:** Must demonstrate understanding of recommendation drivers
- **GDPR Article 22:** Right to explanation for automated decisions
- **FCA/SEC:** Model governance expectations intensifying

Regulators (and clients) demand explanations, not just backtests.

The Traditional Trade-off

The Conventional Dilemma:

You must choose between performance and interpretability

Simple Models

- ✓ Interpretable
- ✓ Regulatory friendly
- ✗ Limited power
- ✗ Miss patterns

Linear regression, logistic regression

The Traditional Trade-off

The Conventional Dilemma:

You must choose between performance and interpretability

Simple Models

- ✓ Interpretable
- ✓ Regulatory friendly
- ✗ Limited power
- ✗ Miss patterns

Linear regression, logistic regression

Complex Models

- ✓ Superior performance
- ✓ Capture nonlinearities
- ✗ Black box
- ✗ Can't explain

Neural networks, XGBoost, random forests

The Traditional Trade-off

The Conventional Dilemma:

You must choose between performance and interpretability

Simple Models

- ✓ Interpretable
- ✓ Regulatory friendly
- ✗ Limited power
- ✗ Miss patterns

Linear regression, logistic regression

Complex Models

- ✓ Superior performance
- ✓ Capture nonlinearities
- ✗ Black box
- ✗ Can't explain

Neural networks, XGBoost, random forests

This is a false choice.

A Bird's Eye View on Explainable AI

Explainable AI (XAI)

What is XAI?

- Post-hoc techniques that explain complex ML models
- Provides explanations without changing the model
- Model-agnostic: works with any algorithm

Explainable AI (XAI)

What is XAI?

- Post-hoc techniques that explain complex ML models
- Provides explanations without changing the model
- Model-agnostic: works with any algorithm

What you get:

- Keep the predictive power of complex models
- Add layer of explanation on top
- Navigate compliance faster and defend decisions to regulators and clients

SHAP: Unpacking Model Predictions

What SHAP does:

- Decomposes a prediction into feature-level contributions
- Shows clearly where actionable insights come from
- Production-ready Python libraries (scikit-learn, XGBoost, TensorFlow, etc).

Benefits for investment professionals:

- Explain decisions to compliance
- Identify which factors drive your decision-making process

Note: SHAP can be computationally heavy for large models.

SHAP Theoretical Foundation

Shapley values from cooperative game theory:

- Features are "players" contributing to a prediction
- Each feature gets credit for its marginal contribution

SHAP Theoretical Foundation

Shapley values from cooperative game theory:

- Features are "players" contributing to a prediction
- Each feature gets credit for its marginal contribution

Key properties for model explainability:

- **Local accuracy:** Feature contributions sum exactly to the prediction (no unexplained remainder)
- **Consistency:** If a feature becomes more important to the model, its SHAP values increase correspondingly
- **Missingness:** Irrelevant features get zero attribution

SHAP in Action: Stock Return Prediction

Analyst: "Stock X will perform +1.6% next month!"

SHAP decomposition explains why:

$$\text{Prediction} = \text{Base Value} + \text{Feature Contributions}$$

$$1.60\% = 0.20\% + 1.40\%$$

Feature contributions breakdown:

+0.2% **Base value** - Historical average return

+1.2% **Momentum** - 12-month price trend

+0.6% **Low volatility** - Stable returns profile

−0.4% **Valuation** - High P/E ratio (overvaluation risk)

SHAP in Action: Stock Return Prediction

Analyst: "Stock X will perform +1.6% next month!"

SHAP decomposition explains why:

$$\text{Prediction} = \text{Base Value} + \text{Feature Contributions}$$

$$1.60\% = 0.20\% + 1.40\%$$

What you can tell your CIO: *"We're long Stock X primarily for momentum exposure. The high valuation is a known risk we're accepting for the stable return profile."*

SHAP in Action: Stock Return Prediction

Analyst: "Stock X will perform +1.6% next month!"

SHAP decomposition explains why:

$$\text{Prediction} = \text{Base Value} + \text{Feature Contributions}$$

$$1.60\% = 0.20\% + 1.40\%$$

What you can tell your CIO: *"We're long Stock X primarily for momentum exposure. The high valuation is a known risk we're accepting for the stable return profile."*

Research example: Bianchi et al. (2025) *"Transaction Costs and the Stochastic Discount Factor."* Working Paper

- `DeepSHAP` to show transaction cost-aware tangency portfolios rely less on high-turnover characteristics

Surrogate Models: Decomposing Model Behavior

What surrogate models do:

- Train a simple model to approximate your complex model's "global" behavior
- Extract interpretable insights from the approximation

The approach:

- Complex model makes predictions on representative data
- Simple model learns to mimic those predictions based on the same input

Rationale: Use original complex model for investing, but use simpler surrogate for explaining.

Surrogate Models in Action

Original model: Deep Learning with 100+ features predicting Stock X returns

Surrogate: Train a decision tree to mimic DL predictions

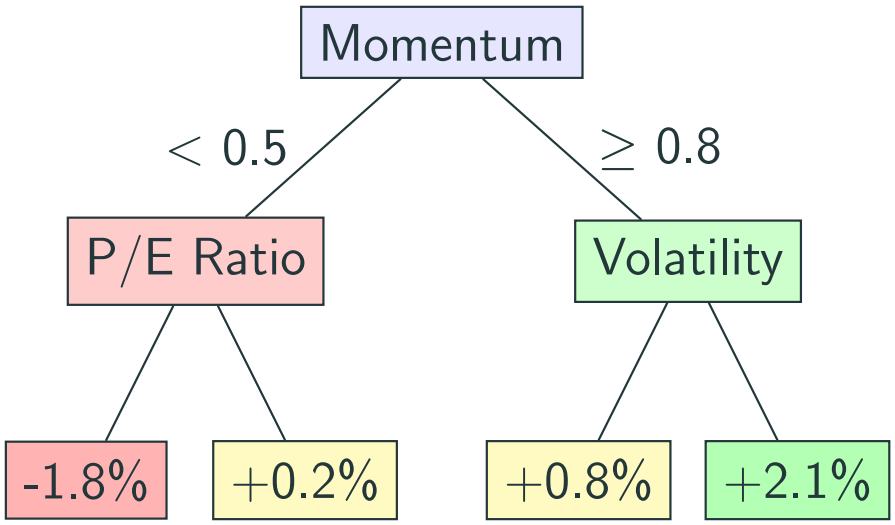

- **If** Momentum ≥ 0.8 & volatility < 0.15
  Then +2.1% prediction
- **If** Momentum < 0.5 & P/E > 25
  Then -1.8% prediction

The tree reveals the DL pattern:

High momentum + low vol $\rightarrow$ strong positive returns.

SHAP vs Surrogate Models: When to Use Each

SHAP Values

- Individual predictions
- "Why this recommendation?"
- Day-to-day operations
- Client reporting

Example:

"We're long Stock X for +1.2%
momentum exposure, despite
-0.4% valuation headwind"

SHAP vs Surrogate Models: When to Use Each

SHAP Values

- Individual predictions
- "Why this recommendation?"
- Day-to-day operations
- Client reporting

Example:

"We're long Stock X for +1.2% momentum exposure, despite -0.4% valuation headwind"

Surrogate Models

- Overall model behavior
- "How does this work?"
- Regulatory docs
- Model validation

Example:

"Model systematically buys stocks with momentum > 0.8 and vol < 0.15"

Main Takeaways

Wrap Up

Two key points:

1. Explainability is now mandatory, not optional
2. XAI lets you keep complex models while adding transparency

Practical next steps:

- Audit your models: which lack explanations?
- Start with SHAP and document your approach for compliance and validation

Bottom line: You can deploy complex models for investment decision-making without sacrificing transparency.

Thank You

Questions?

Contact:

d.bianchi@qmul.ac.uk

Slides and research available at:

whitesphd.com

Selected Research: XAI Applications in Finance

Bianchi, D., Jiao, T., and Ma. H. (2025). "Transaction Costs and the Stochastic Discount Factor." *Working Paper*

- Uses DeepSHAP to reveal how transaction costs reshape characteristic importance

Bianchi, D. and Jiao, T. (2024). "Macroeconomic Fundamentals and the Shape of Sovereign Credit Risk." *Working Paper*

- Uses SHAP + Partial Dependence Plots to identify macro drivers