Module 5 Summary: The Automation Problem

Prof. Dr. Joerg Osterrieder

Digital Finance — BSc Course

Module 5: Prerequisites

Theme: Automation (ML foundations, GenAI/LLMs, limits, MLOps) **From prior modules:**

- **M2L4:** algorithmic-fairness vocabulary
- **M4L1:** market data and time-series intuition

External knowledge assumed:

- Basic linear algebra (vectors, matrices, dot product); one undergraduate course
- Calculus intuition (gradient, chain rule) at the level of one paragraph, not derivation
- Python with NumPy and scikit-learn at the level of fitting a simple model

Will be introduced this module: We introduce supervised-learning foundations, GenAI/LLMs for finance, the limits of prediction (Goodhart, base-rate neglect, regime shift), and MLOps for production deployment.

Prerequisites are advisory; lessons remain self-contained where feasible. Forward references inside lessons flag any concept used before its canonical introduction.

Key Concepts by Lesson

L1: ML Foundations

- Classification ("what kind?") vs. regression ("how much?")
- Bias-variance trade-off: underfitting vs. overfitting
- Train/validation/test split; cross-validation
- Feature engineering for financial time series

L2: Generative AI & LLMs

- Transformer architecture and attention mechanism
- LLMs: next-token prediction, not "understanding"
- Financial use cases: document processing, compliance review
- RAG pipelines for financial document Q&A

L3: Limits of Prediction

- Stationarity: mean, variance, autocorrelation must be constant
- Regime changes break models trained on historical data
- Pitfalls: data snooping, look-ahead bias, survivorship bias

L4: MLOps

- MLOps = DevOps principles for ML lifecycle
- Drift types: concept drift, data drift, covariate shift
- Model monitoring, versioning, and governance
- Production gap: notebook $\neq$ production system

Module 5 answers: Can machines make better financial decisions than humans — and at what cost?

Key Formulas and Definitions

Bias-Variance Decomposition

$$\text{Expected Error} = \text{Bias}^2 + \text{Variance} + \text{Irreducible Noise}$$

High bias = underfitting (model too simple). High variance = overfitting (model too complex).

Stationarity Requirement

$$E[y_t] = \mu, \quad \text{Var}(y_t) = \sigma^2, \quad \text{Cov}(y_t, y_{t+k}) = \gamma(k) \quad \forall\, t$$

Financial data almost always violates stationarity — regimes shift, volatility clusters, trends emerge.

Overfitting Diagnostic

If training error $\ll$ validation error $\rightarrow$ overfitting. If both are high $\rightarrow$ underfitting. If converging at low level $\rightarrow$ good fit. Plot learning curves (error vs. data size) to diagnose.

Drift Detection

Concept drift: $P(Y|X)$ changes. **Data drift:** $P(X)$ changes. **Covariate shift:** $P(X)$ changes but $P(Y|X)$ stays the same. All require retraining or model replacement.

The hardest part of financial ML is not building the model — it is knowing when the model has silently stopped working.

Connections to Other Modules

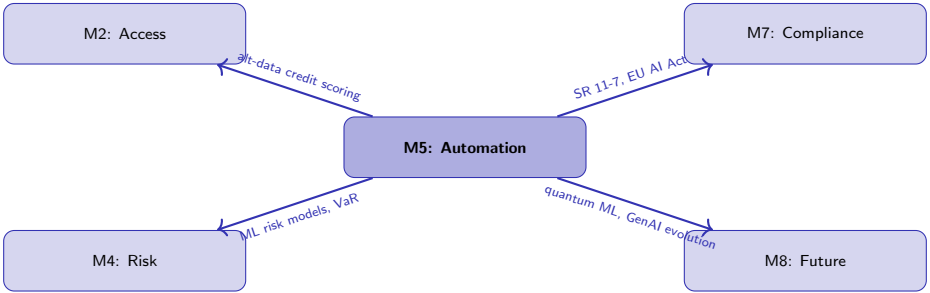

- **Automation → Access (M2):** ML powers alternative credit scoring that includes the "credit invisible" — but introduces algorithmic bias
- **Automation → Risk (M4):** ML models estimate VaR, detect fraud, and predict default — model risk is the cost of automation
- **Automation → Compliance (M7):** SR 11-7 governs model risk; the EU AI Act classifies credit scoring as "high risk" requiring human oversight
- **Automation → Future (M8):** Quantum computing may accelerate ML training; GenAI capabilities will continue to expand

ML is the engine; data is the fuel; governance is the brake. All three must work together.

Module 5: Integrative Questions

Two questions that need more than one lesson to answer:

1. M5L2 (GenAI/LLM) and M5L3 (limits of prediction) suggest different things about reliability. Sketch a concrete finance task where M5L2 succeeds AND M5L3 says you should still distrust the answer.
2. Combine M5L1 (ML foundations) + M5L4 (MLOps) to explain why a model with 95 percent backtest accuracy can degrade to 60 percent in three months of production drift.

Use these as study prompts before the module exam; each integrates concepts that span lessons.

Module 5: Worked Multi-Lesson Example

Worked example: deploy a transaction-classification model from M5L1 (gradient-boosted trees on labelled data), wrap with M5L2 (LLM-assisted feature engineering), check against M5L3 (regime shift detection), operate via M5L4 (CI / model registry / shadow deployment). **Pedagogical pattern:** the example is intentionally end-to-end. Solve it lesson-by-lesson, then step back and identify the lesson whose assumption was the binding constraint.

The exam-style version of this example appears in extttv4/exam_prep/exam_bank.tex for module 5.

Module 5: Forward-Pointing Dependency Map

Concepts from Module 5 that later modules will use:

- **M7L4:** M5L4 MLOps is the day-to-day instantiation of M7L4 model-risk governance
- **M7L3:** EU AI Act compliance (M7L3) operates over the artefacts produced by M5L1-L4
- **M8L2:** Quantum-readiness (M8L2) requires retraining the cryptography assumptions inside M5L1 feature pipelines

Forward-pointing dependencies; concepts not in this map are local to Module 5.