POS Lab: Algorithmic Trading

Dr. Yves J. Hilpisch[1]

December 11, 2025

Purpose and Connection

This lab invites you to apply ideas from *Introduction to the Philosophy of Science: Concepts, Practice, and Case Studies* [1] and from Timmermann and Granger's survey on the efficient market hypothesis and forecasting [2] to the design and critique of algorithmic trading strategies. You will explore randomness, empirical support for and against the EMH, nonlinearity, high dimensionality, and the fragility of statistical and economic "edges" in real markets.

[1] Get in touch: `https://linktr.ee/dyjh`. Web page `https://hilpisch.com`. Research, structuring, drafting, and visualizations were assisted by GPT 5.1 as a co-writing tool under human direction.

Contents

1 Scenario: An Alpha-Seeking Trading Desk

Imagine you are part of a quantitative trading team at a mid-sized hedge fund. Your mandate is to design alpha-seeking strategies on liquid futures and equities, aiming to outperform a risk-adjusted benchmark in the presence of many other sophisticated players.

High-Level Narrative

- The desk runs dozens of strategies, from simple trend-following rules to complex machine-learning models.
- Management believes that markets are "mostly efficient" but that small, short-lived inefficiencies can be monetised.
- Regulators and clients care about risk management, robustness, and the possibility that models themselves may affect market stability.

Your task in this lab is to design and interrogate an algorithmic trading research pipeline using the philosophical and statistical tools from the main book [1] and the forecasting perspective of Timmermann and Granger [2].

From Lab to Life

Keep at least one concrete asset class in mind (for example equity index futures, FX, or liquid single-name stocks). As you work through the sections, imagine writing an internal research memo that would convince a sceptical risk committee that your strategy is:

- empirically supported rather than overfit,
- computationally and operationally feasible,
- robust to modest changes in market conditions.

2 Data, Measurement, and Market Microstructure

Relate this section to Chapter 5 on measurement and error and Chapter 13 on quantitative finance in the main book [1].

Step 1: Defining the Target Quantity

- Decide what your primary target is: excess return over a benchmark, Sharpe ratio, risk-adjusted profit after costs, or something else.
- Write the target as a symbolic expression (for example daily strategy return R_t^{strat} relative to market return R_t^{mkt}).
- List the data sources needed to compute this target: prices, volumes, corporate actions, transaction costs, risk-free rate proxies.

Guiding Questions

- Where do survivorship bias and look-ahead bias creep in when you construct your dataset?
- How do bid–ask spreads, slippage, and latency distort the apparent edge you see in cleaned historical data?
- Which pre-processing steps (resampling, de-meaning, volatility scaling) risk smuggling in information from the future?

Step 2: Recognising Noise and Microstructure Effects

Sketch, conceptually or with a quick plot, how observed returns at very short horizons differ from "true" underlying price changes:

- At high frequency, quote changes may largely reflect bid–ask bounce and order-book noise.
- At daily or weekly horizons, the same series may look smoother and more amenable to modelling.
- Structural breaks (for example tick-size changes, fee structure changes) can alter microstructure noise patterns.

Mark which horizons you consider meaningful for your alpha idea and which are dominated by microstructure effects you would rather treat as noise.

3 Forecasting and the Efficient Market Hypothesis

Relate this section to Chapter 2 on evidence, Chapter 7 on overfitting, and the survey by Timmermann and Granger [2].

Step 3: Linear Forecasts and EMH Flavours

Consider a simple forecasting regression

$$R_{t+1} = \alpha + \beta^\top X_t + \varepsilon_{t+1},$$

where R_{t+1} is next-period excess return and X_t is a vector of predictors (for example lagged returns, valuation ratios, macro indicators).

- Classify which EMH flavour (weak, semi-strong, strong) would be challenged if a strategy built from in-sample estimates (using $\hat{\beta}$) delivered reliably non-zero abnormal returns out-of-sample.
- Decide how you would split data into estimation and evaluation windows, following the spirit of forecast evaluation in [2].

- Specify at least two loss functions for assessing forecasts (for example squared error vs. utility-weighted loss).

In a short paragraph, explain what kind of evidence—effect size, stability across subsamples, economic significance after costs—would make you comfortable calling this a genuine predictive relation rather than noise.

Step 4: Multiple Testing and Data-Snooping

List at least five researcher degrees of freedom in your forecast specification:

- choice of look-back window and sampling frequency,
- inclusion or exclusion of certain predictors,
- transformations (log, differences, volatility scaling),
- treatment of missing data and outliers,
- decision to report or omit particular subsample results.

For each, note how you would constrain or pre-register choices to avoid data-snooping while still allowing legitimate model comparison. Relate this explicitly to the overfitting and p-hacking discussions in the main text [1].

4 Nonlinearity, High Dimensionality, and Computation

Relate this section to Chapters 7, 9, and 12.

Step 5: Nonlinear and Machine-Learning Models

Suppose you replace the linear regression with a nonlinear model (for example gradient boosting, random forests, or a neural network).

- List three ways in which such a model can capture structure that a linear model cannot (for example interactions, thresholds, regime changes).
- List three new ways in which such a model can go wrong (for example extreme sensitivity to hyperparameters, instability across periods, spurious fits to rare regimes).
- Describe your plan for cross-validation or walk-forward testing, including how you would avoid leakage from the test set into training through hyperparameter tuning.

Step 6: Curse of Dimensionality and Compute Budgets

Connect your model choice to the curse of dimensionality and computational limits:

- Estimate, even roughly, how the number of free parameters in your model grows with the number of features and layers.
- State what daily or weekly compute budget you are willing to spend on retraining and evaluation.
- Identify one simplification (for example factorising features, using lower-frequency aggregates, or imposing sparsity) that trades a little expressiveness for a large gain in stability or speed.

Summarise in a few sentences how these constraints influence which models you deem scientifically and commercially acceptable, not just technically possible.

5 Instability, Regimes, and Edge Durability

Relate this section to Chapter 9 on complexity and the finance case study in Chapter 13.

Step 7: Regime Diagnostics

Design a simple regime-detection or diagnostics protocol:

- Choose at least one volatility measure and one cross-sectional dispersion measure to track.
- Define thresholds or change-point rules that you suspect mark different regimes (for example low-vol vs. high-vol, crisis vs. calm).
- Describe how you would record strategy performance by regime and how that would feed back into capital allocation or risk limits.

Explain how you would distinguish between a genuinely broken strategy and a temporary drawdown that is still consistent with the original edge.

Step 8: A Schematic Research Loop

Figure 1 offers a schematic view of an algorithmic trading research and deployment loop. As you look at it, annotate where in your own workflow:

- new data enter,
- models are updated or retired,
- risk and compliance gates constrain deployment,
- feedback from P&L and market impact flows back into your hypotheses.

6 Reflexivity and Model-Aware Markets

Relate this section to the reflexivity discussion in Chapter 13 and the complex-systems themes in Chapter 9 of the main book [1].

Step 9: Self-Fulfilling and Self-Defeating Strategies

List one example strategy that could become self-fulfilling and one that could become self-defeating:

- For the self-fulfilling case, describe a situation where widespread adoption of a rule (for example "buy on 200-day moving-average cross") temporarily makes its own signal profitable.
- For the self-defeating case, describe how crowding into a once-profitable niche (for example a specific ETF arbitrage) can erode or reverse its edge.

Explain how you would monitor for these feedbacks in live trading and which indicators (flows, spreads, liquidity, slippage) would warn you that your own models are part of the regime shift.

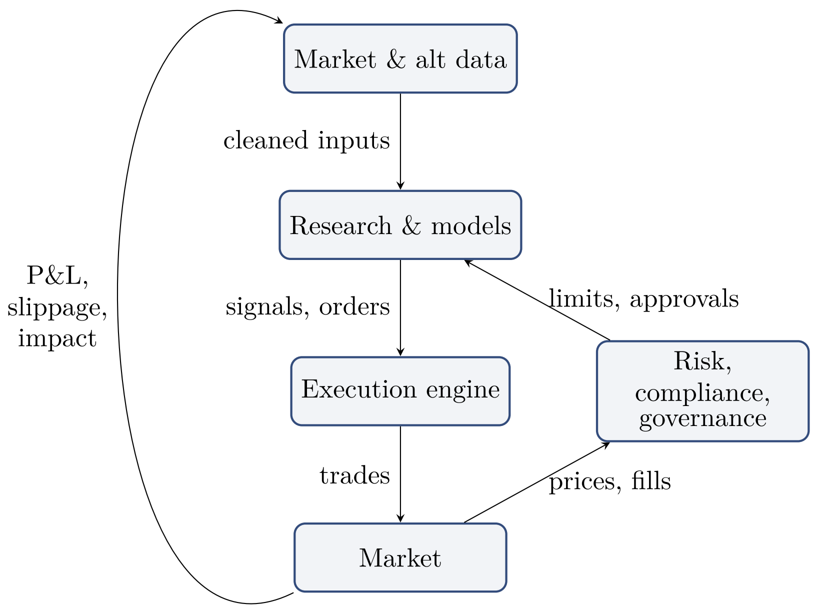

Figure 1: Schematic algorithmic trading loop: data feed research, research generates signals, execution sends orders to the market, and risk/governance plus realised P&L feed back into model design and deployment.

Step 10: Governance and Philosophy-of-Science Checks

Translate key ideas from the book into concrete governance questions for your desk:

- **Measurement**: Are the inputs to your models (signals, risk metrics) still measuring what you think they measure?
- **Causality**: Where do you rely on causal stories (for example risk premia) versus purely statistical correlations?
- **Evidence**: How do you balance in-sample fit, out-of-sample performance, and multi-market replication?
- **Complexity**: Are you using model complexity where it adds genuine value, or mainly because it is fashionable?

Write a short checklist (5–7 bullets) that you would want every new strategy proposal to address before it is approved.

7 Putting It Together: A Mini Strategy Dossier

As a final exercise, draft a one-page internal "strategy dossier" that:

- states the economic and statistical hypotheses behind your alpha,
- summarises the data sources and main measurement issues,
- describes the forecasting model(s) and how they were validated,
- explains how you account for high dimensionality and computational limits,
- reports performance by regime and discusses edge durability,
- identifies key reflexive feedback channels and governance safeguards.

Try in 60 Seconds

If time is short:

- Write down one reason you think markets are *more* efficient than many popular trading books suggest, and one reason you think they are *less* efficient.
- Name one way in which your own modelling choices could make the market a little more efficient, and one way they could accidentally amplify instability.

References

[1] Y. J. Hilpisch. *Introduction to the Philosophy of Science: Concepts, Practice, and Case Studies*. 2025. Available at `https://hilpisch.com/philosophy.pdf`.

[2] A. Timmermann and C. W. J. Granger. Efficient market hypothesis and forecasting. *International Journal of Forecasting*, 20(1):15–27, 2004. Available at `http://www.e-m-h.org/TimmermannGranger2004.pdf`.