30-MINUTE ORIENTATION LECTURE

Quantitative Investment Systems

Data, AI, Portfolio Engineering and Execution

From mathematical models to **defensible investment decisions**

ORIENTATION FOCUS

What the course trains, how students work, and why it matters in AI-era quant finance.

Dr. Don Huang
Co-founder & Chief Scientist, AQUMON
PhD in Mathematics, HKUST
29 August 2026

A 30-minute roadmap

Central question for the semester

How do we turn **noisy information** into a **defensible investment decision**?

PART 1

Course logic

decision system

PART 2

Learning setup

backbone + experiments

PART 3

Career lens

evidence + defense

Build a model you can trust, implement, and defend.

The course starts where an idealized model ends

Textbook abstraction

known returns

stable covariance

optimal portfolio

frictionless trading

Investment reality

estimated inputs

changing regimes

noise amplification

costs and frictions

Every quantitative decision should be tested, challenged, and monitored.

Quantitative investing as a decision system

Information → **Model** → **Forecast** → **Portfolio**

Trade → **Monitoring** → **Revision / Retirement**

Research risk
timing, leakage, stability

Decision risk
estimates, constraints, sensitivity

Implementation risk
cost, liquidity, stop rules

A beautiful backtest

ARNOTT, HARVEY AND MARKOWITZ (2019): “ALPHA-BET” BACKTEST

It first looks real

- NYSE market-neutral: 1963–2015
- Built on 1963–1988; works in 1989–2015
- Nearly 6% alpha; low turnover

What would you question first?

Long–Short Market-Neutral Strategy Based on NYSE Stocks, January 1963 to December 2015

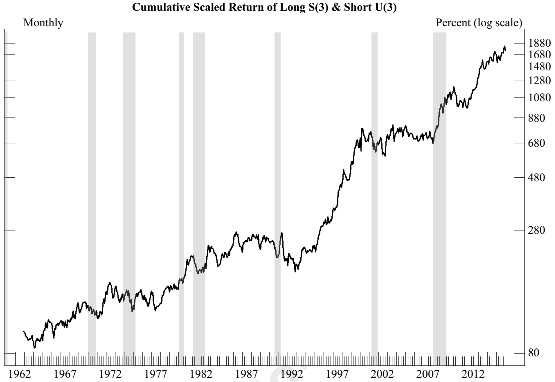

Notes: Gray areas denote NBER recessions. Strategy returns scaled to match S&P 500 T-bill volatility during this period.
Source: Campbell Harvey, using data from CRSP.

ARNOTT, HARVEY AND MARKOWITZ (2019): “ALPHA-BET” BACKTEST

It first looks real

- NYSE market-neutral: 1963–2015
- Built on 1963–1988; works in 1989–2015
- Nearly 6% alpha; low turnover

Punchline

Long third-letter S; short third-letter U.

No real economic story.

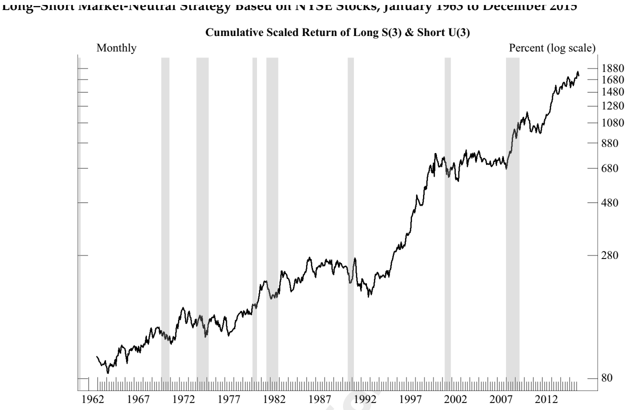

Notes: Gray areas denote NBER recessions. Strategy returns scaled to match S&P 500 T-bill volatility during this period.
Source: Campbell Harvey, using data from CRSP.

The 13-week journey

Unifying thread

How does uncertainty propagate from data to portfolio outcome?

Foundations $\longrightarrow$ **Research** $\longrightarrow$ **Portfolio**

AI & Trading $\longrightarrow$ **Defense**

Point-in-time data | honest backtesting | portfolio construction | LLMs | project defense

What you will actually do

Research layer

- point-in-time data
- risk and covariance models
- alpha diagnostics
- flawed-backtest audit

Portfolio and execution layer

- constrained portfolios
- direct indexing / tax-aware choices
- ML and LLM signals
- execution and cost

Python is the implementation language. Research judgment is the subject.

The coding requirement: enough Python to experiment

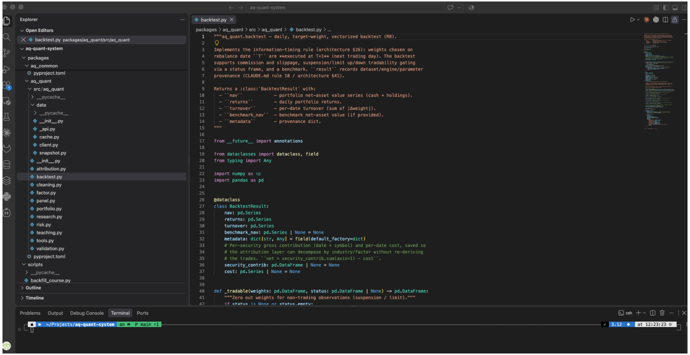

REQUIRED	PROVIDED	FOCUS
working Python	**quant engine backbone**	**research judgment**

Why this belongs in Financial Mathematics

The bridge this course tries to build

Models $\longrightarrow$ **Inputs** $\longrightarrow$ **Decisions** $\longrightarrow$ **Validation**

What you bring

probability, statistics, linear algebra, optimization, financial modeling

What the course adds

data discipline, implementation frictions, validation, monitoring, defense

Turn mathematical ideas into defensible financial decisions.

What I hope you can take away

Research Integrity
no leakage; no backtest luck

Decision Engineering
forecasts under risk, turnover, cost

Model Defense
assumptions, failure modes, monitoring

The final product is not "a good backtest."

It is a defensible investment decision process.

The semester project: build it, then defend it

Build

problem $\rightarrow$ data $\rightarrow$ forecast $\rightarrow$ portfolio $\rightarrow$ implementation $\rightarrow$ validation

Defend

- most uncertain input
- drivers of holdings
- simpler benchmark
- cost survival
- stop-trading rule

Expected background: probability/statistics, linear algebra, basic investments, and working Python. Convex optimization helps.

PART II

How can the course help you become useful in a quant finance team?

Part I explains what you build. Part II explains why that evidence matters.

Working hypothesis

AI makes implementation easier; this course builds the **judgment** to decide when a model is usable.

AI changes the career equation

Becoming cheaper

- first prototypes
- boilerplate code
- baseline models
- text summarization

Becoming more valuable

- right question
- leakage control
- portfolio usefulness
- risk communication

Use AI for implementation. Invest your learning time in decision judgment.

Course work becomes career evidence

Course work

- point-in-time data
- alpha validation
- portfolio construction
- ML / LLM evaluation
- project defense

Career signal

- research discipline
- statistical honesty
- decision awareness
- AI-finance literacy
- communication under scrutiny

Quant Research | Portfolio Engineering | AI / ML for Finance | Model Validation

Make the semester project interview-ready

A strong project tells one coherent story

hypothesis $\rightarrow$ data $\rightarrow$ validation $\rightarrow$ portfolio decision $\rightarrow$ risk control

Show

economic rationale, data timing, robustness, cost impact

Defend

fragile assumptions, benchmark, monitoring, retirement

The goal is not the fanciest algorithm; it is a defensible decision process.

Recap: what to take away

1

Think in systems

2

Practice judgment

3

Create evidence

If there is one phrase to remember:

Build something you can defend.

Questions and discussion