

Quantitative Finance

HEC Fintech Summer School

Edoardo Vittori

July 7th, 2026

Can We Predict the Future?

1

Every trade is a bet on the future

Markets are hard to predict, but not impossible

Example: EUR/USD, at every moment we choose to buy, sell, or wait

Weather is partly predictable. Markets are harder: they are competitive, they adapt, and most of what moves them is noise. We will not look for certainty — we will look for a small, repeatable edge.



BUY

You expect EUR/USD to rise



WAIT

No clear edge right now



SELL

You expect EUR/USD to fall

Now: 1.1385. The only question that pays is what to do next.

A small edge, repeated, can pay

One coin flip is luck. Ten thousand flips with a 51% edge is a business

Bet €1 that EUR/USD ticks up. With a 51% chance of being right, you barely beat a coin toss on any single trade. But independent, equally-sized bets accumulate

	50% — no edge	51% — tiny edge
Edge per €1 bet	€0.00	+€0.02
After 10,000 bets	~€0 (pure luck)	+€200 expected
Chance you finish with a profit	~50% (coin toss)	>97%

The edge is real only if the bets are independent and identically sized. Break that, and the math breaks.

Costs, asymmetry and risk break the fair game

Three forces quietly turn a winning coin-flip into a losing business

1 · Costs

+0.020 − 0.018 = 0.002

Every trade pays the spread, commission and impact. Trade often and costs can swallow the edge.

2 · Asymmetry

99 × +€1 − 1 × €200

Many small wins and one big loss. A high hit rate can still lose money.

3 · Risk & drawdown

−50% needs +100%

Losses compound against you. Survival, not the next trade, comes first.

Prediction is not trading

A forecast tells you what might happen; a decision tells you what to do about it

PREDICTION asks

“Will EUR/USD go up or down?”

- Output is one number: a probability
- A forecasting problem
- Graded by accuracy
- Says nothing about how much to bet

TRADING asks

“Given probability, payoff, costs, risk, do I buy, sell, or wait and what size?”

- Output is an action and a size
- A decision problem under uncertainty
- Graded by P&L after costs and risk
- This is what actually makes money

A good prediction is an input. A good decision is the output.

Where we're going today

1	**Can we predict the future?**	Prediction vs trading
2	**Capital Markets landscape**	Who trades, why, and where you fit
3	**Financial engineering foundations**	Models for uncertainty, pricing & execution
4	**Systematic strategies**	From market intuition to backtested P&L
5	**Portfolio optimization**	Combining many imperfect strategies
6	**Machine learning in trading**	Extending the workflow — without the hype
7	**Gen AI on equity market neutral**	Using Gen AI on single stocks

Capital markets players

Player	The decision they care about
Corporate Client	How do I hedge my FX / rate / commodity risk?
Sales	Which solution fits this client?
Structurer	What payoff solves the client's problem?
Market maker	What bid/ask spread and inventory do I run?
Quant	What model, signal or tool supports the decision?
Risk manager	What can go wrong, and how much?
Asset manager	How do I allocate across assets over time?
Hedge fund / prop	Can this signal become a repeatable strategy?
Execution Trader	How do I execute and reduce transaction costs?

Market making

Offering liquidity to the markets

Regulated market example

Last	Last Vol	Total Vol	Close	Daily Low	Daily High
4045.00	2	367267	4097.50	4033.50	4101.50

Implied

Bid		Offer	
Volume	Price	Price	Volume
136	4044.50	4045.00	62
327	4044.00	4045.50	293
348	4043.50	4046.00	427
620	4043.00	4046.50	426
358	4042.50	4047.00	463
330	4042.00	4047.50	348
325	4041.50	4048.00	327
318	4041.00	4048.50	294
305	4040.50	4049.00	281
512	4040.00	4049.50	288

Dealer market example - OTC

MARKIT ITRX EUR SNR FIN 06/26 | 92) Order Book | 93) RFS | 97) Settings

11:56:39 | 95) Buy | 96) Sell | BTFE | Filter By | All

PCS	Firm Name	CCP	Bid Spd	Ask Spd	BSz(MM)	ASz(MM)
CSDE	CREDIT SUISSE INTL	ICEE	54.6900 / 55.0100		50 x 50	
CCGC	Citi CCGC	ICEE	54.7650 / 55.0350		50 x 50	
GSMX	GS MINI	ICEE	54.7350 / 55.0350		15 x 15	
JCTT	JP MORGAN	ICEE	54.7600 / 55.0400		100 x 100	
BXCZ	Barclays Minis	ICEE	54.8400 / 55.0400		75 x 75	
MSTI	MORGAN STANLEY MINI	ICEE	54.8000 / 55.0400		50 x 50	
ABNP	BNP Paribas	ICEE	54.8000 / 55.0500		51 x 51	
SGMI	SocGen Mini	ICEE	54.7380 / 55.0880		50 x 50	
CSEO	CS iTraxx Mini	ICEE	54.610 / 55.090		100 x 100	
BARX	Barclays	ICEE	54.7650 / 55.1150		250 x 250	
CGCX	Citi CGCX	ICEE	54.6800 / 55.1200		100 x 100	
EBNP	BNP Paribas	ICEE	54.7250 / 55.1250		101 x 101	
SCDS	SocGen	ICEE	54.6890 / 55.1380		125 x 125	
DBVD	DB Index-(DBDV)	ICEE	54.8500 / 55.1500		100 x 100	
GSET	GOLDMAN SACHS	ICEE	54.6100 / 55.2100		75 x 75	
CCGB	Citi CCGB	ICEE	54.5600 / 55.2400		200 x 200	
JPOS	JP Morgan	ICEE	54.5600 / 55.2400		200 x 200	
MSTT	MORGAN STANLEY MAXI	ICEE	54.5500 / 55.2900		100 x 100	
CSXE	Credit Suisse EU	ICEE	54.406 / 55.294		200 x 200	

RFQ Example

Client buys protection 200mln
Price: ____

Send

Asset Management and Hedge Funds

Portfolio managers

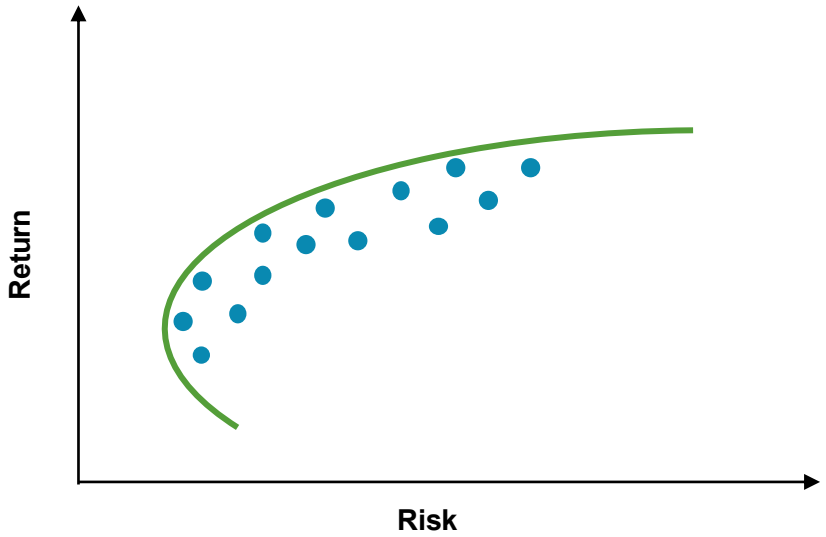
Objective

Manage investments to meet specific risk-return objectives for clients or proprietary funds. Usually comparing to a **benchmark**

Core Activities

- **Portfolio Construction**: apply techniques like Modern Portfolio, Risk Parity, and to construct diversified portfolios.
- **Strategy Allocation**: allocate capital across asset classes and strategies (e.g. long/short equity, macro, credit, multi-strategy) to:
  - Achieve a balanced exposure
  - Maximize risk-adjusted returns (Sharpe ratio, IR)
  - Maintain diversification and liquidity
- **Risk Management**: monitor volatility, VaR, drawdowns, and correlation

Illustration of efficient frontier



Execution desks

Given a trade, execution desks should minimize transaction costs and market impact

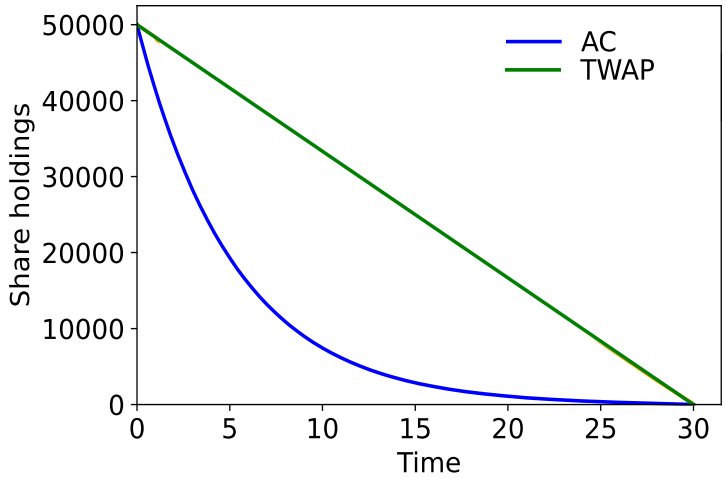
Definition of optimal execution

Given a trade to execute in a specified amount of time, minimize market impact and transaction costs

Order types

- **Market order** is an order to execute immediately at the best price possible
- **Limit order** is an order that specifies both the price and volume of a trade
- A limit order sits in the order book until it is either executed against a matching market order or canceled

Illustration of execution trajectories



Prop Trading

Traders

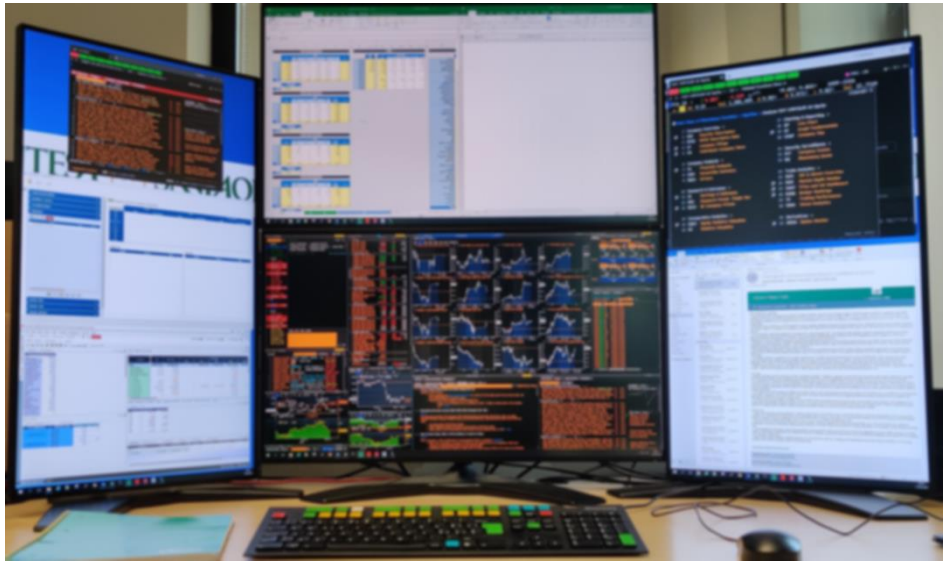
Objective

Generate alpha by exploiting short to medium term market opportunities using the institution's own capital, within strict risk constraints (e.g., VaR limits)

Core Activities

- Alpha Generation: Develop and execute strategies across asset classes (e.g., equities, FX, rates, commodities, crypto) based on:
  - Technical indicators
  - Quantitative signals (statistical arbitrage, momentum)
  - Macro views or event-driven insights
- **Risk Budgeting:** operate within daily and aggregate Value-at-Risk (VaR) constraints
- **Monitoring:** monitor intraday P&L, position limits, and tail risk

Trading Working Desk Example



Quants and Financial engineering

The enabling component

Quants

Quantitative analysts/researchers/developers use advanced mathematics, statistics, and programming to **analyze financial data** and **develop models** that support trading, risk management, and investment decisions

Roles:

- Data analysis & signal generation for trading strategies
- Backtesting models and calibrating parameters
- Supporting portfolio optimization and asset allocation
- Implementing models in code (usually Python)

Financial Engineering

Application of mathematical models and financial theory to design and **price complex financial instruments** and develop innovative financial products

Roles:

- Quantitative modeling
- Risk management
- Derivatives pricing
- Designing execution algorithms
- Implementing pricing models in code (usually C++)

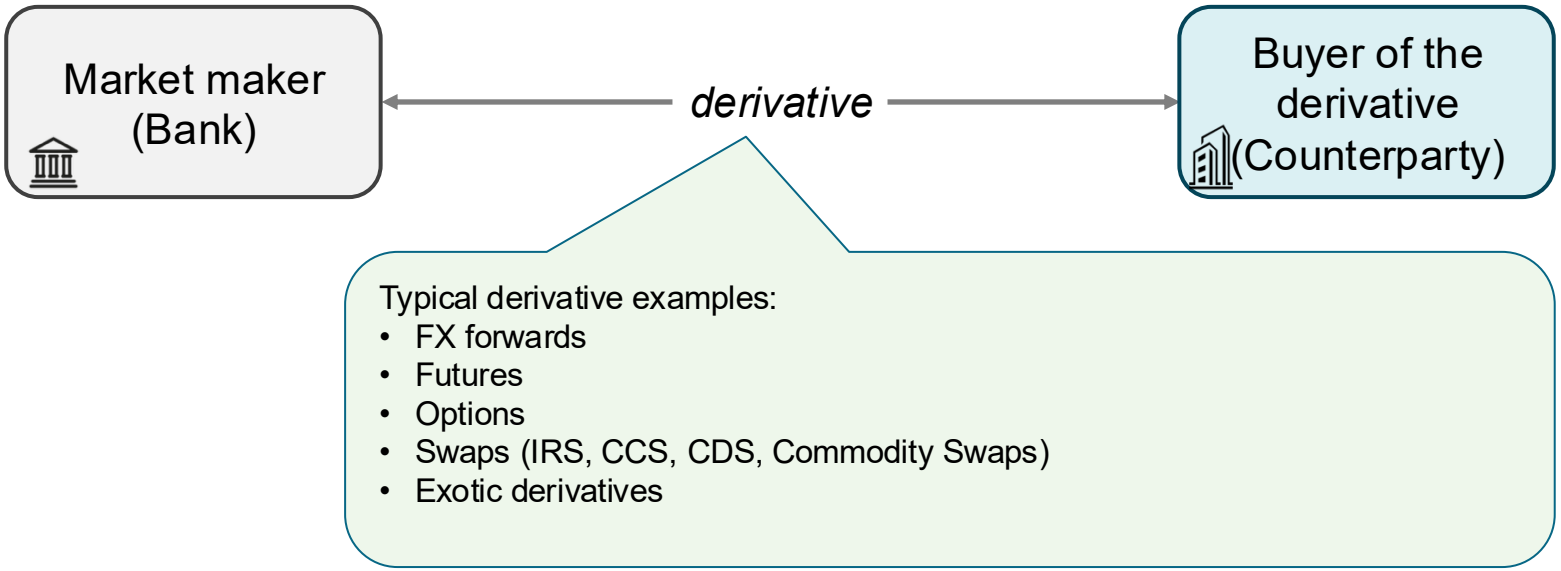
Asset classes and instruments

Each asset class has a range of instruments which can be used

Asset class	Instruments
Equities	Stocks, ETFs, derivatives
Fixed Income (rates)	Government bonds, ETFs, derivatives
Credit	Corporate bonds, ETFs, derivatives
FX	Spot, ETFs, derivatives
Commodities	ETFs, derivatives
Crypto	Spot, ETFs, derivatives

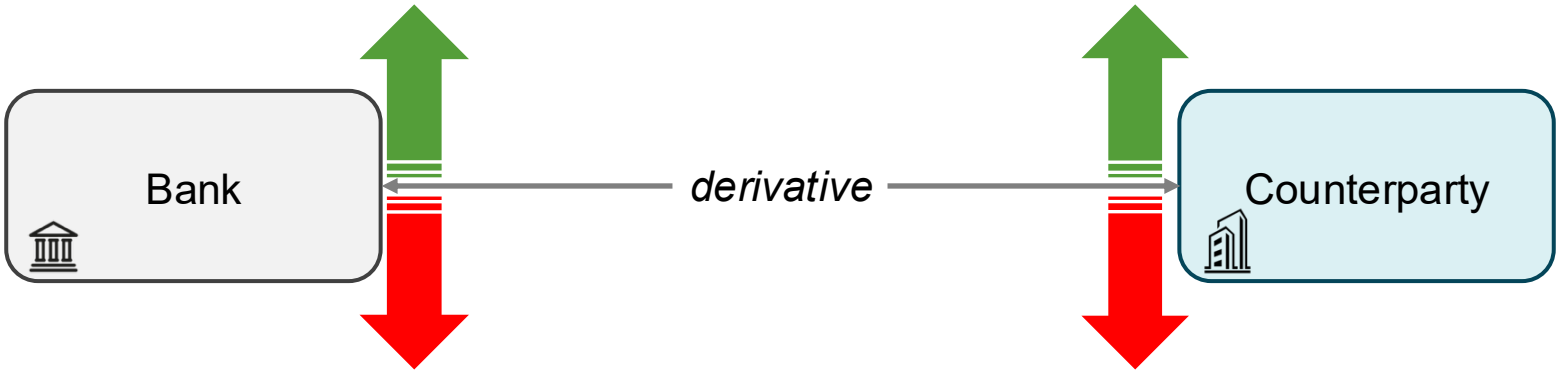
Derivatives

Financial contracts whose value is derived from an underling asset



Derivatives

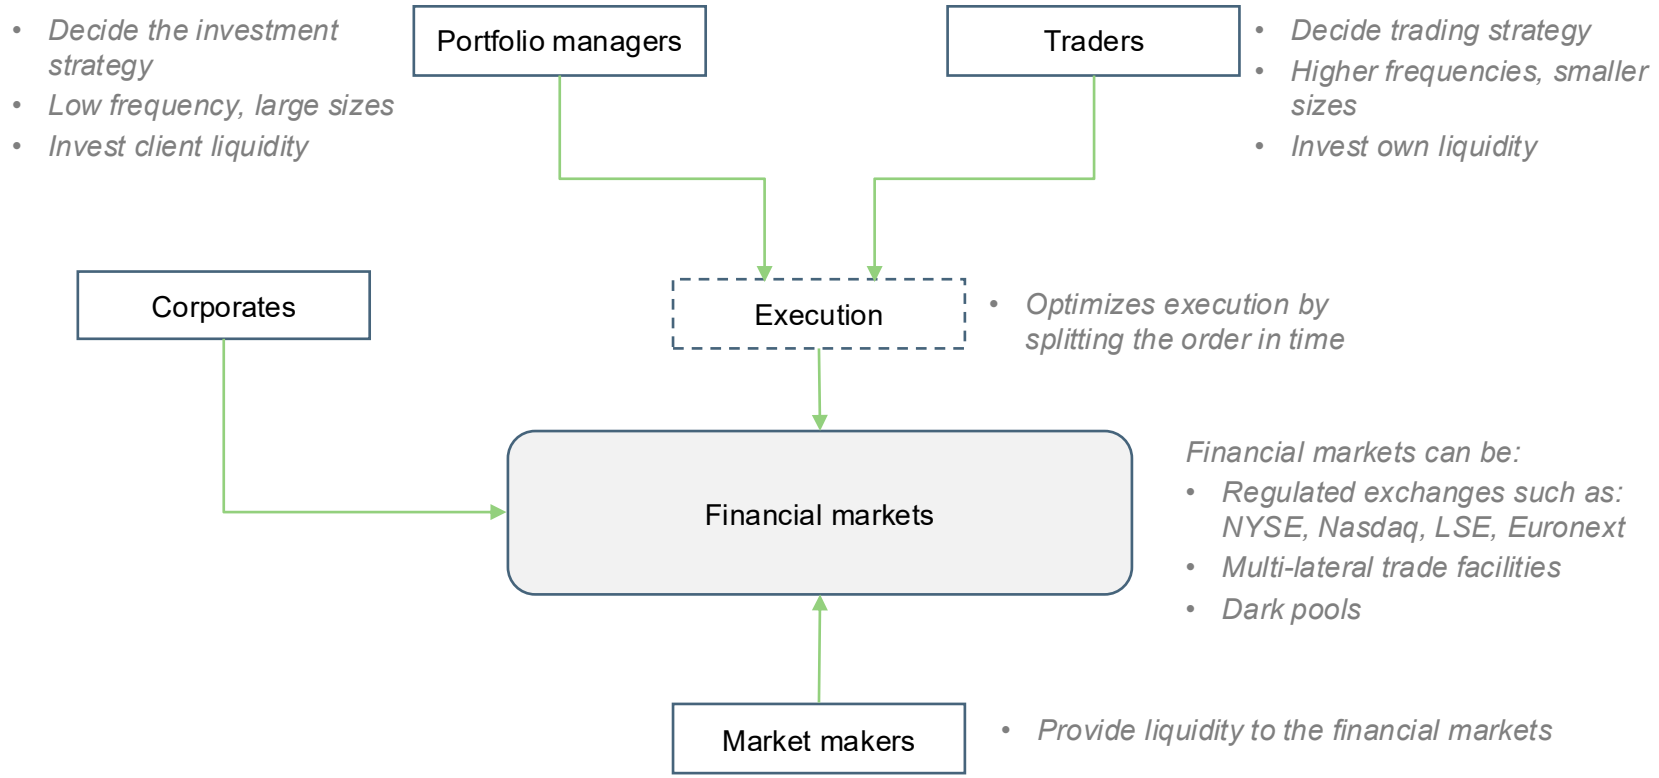
The derivative initially has a zero value, then its value fluctuates based on market conditions



Hedging derivatives
Banks do not like being exposed to unpredictable market movements, so they **hedge** their position by opening an opposite trade with another bank or hedging with correlated but less expensive instruments

Overview of Financial Markets

Focus on the most influential actors



Where we're going today

1	**Can we predict the future?**	Prediction vs trading
2	**Capital Markets landscape**	Who trades, why, and where you fit
3	**Financial engineering foundations**	Models for uncertainty, pricing & execution
4	**Systematic strategies**	From market intuition to backtested P&L
5	**Portfolio optimization**	Combining many imperfect strategies
6	**Machine learning in trading**	Extending the workflow — without the hype
7	**Gen AI on equity market neutral**	Using Gen AI on single stocks

Five tools, five problems

We introduce each concept only after the problem it answers

Model	The problem it solves
Brownian motion	Prices are uncertain and move over time — how do we describe **random price paths**?
Black–Scholes	What is a **fair price for an option?**
Monte Carlo	The future has **many possible paths** — how do we handle them all at once?
Almgren–Chriss	Trading itself moves the market — **how fast should we trade**?
Avellaneda–Stoikov	How should a market maker **quote bid and ask** under inventory risk?

Deterministic or stochastic?

Weather models vs financial markets — same toolbox, opposite treatment of uncertainty

Weather — deterministic PDEs

- Physics fixes the dynamics: Navier-Stokes & thermodynamics evolve today's atmosphere forward
- The equations contain no randomness — same input, same forecast, every run
- Uncertainty comes from imperfect initial conditions, amplified by chaos
- Fix: run ensembles — many forecasts from slightly perturbed starting points

vs

Markets — stochastic SDEs

- No physical law drives prices — millions of adaptive agents do
- Randomness is built into the model itself: $dS = \mu S\,dt + \sigma S\,dW$
- The dW term injects noise at every instant — we model the statistics of shocks, not their cause
- Same fix: Monte Carlo simulates thousands of possible price paths

The twist: Itô's lemma turns the stochastic price model back into a deterministic PDE for the option price — Black-Scholes.

Introduction to stochastic processes

Modeling uncertainty in finance

A **stochastic process** is a collection of random variables $\{X_t\}_{t>0}$ indexed by time.

Brownian motion $\{W_t\}_{t>0}$ is a key stochastic process used in finance:

- $W_0 = 0$ and $W_{t+\Delta t} - W_t \sim N(0, \Delta t)$
- W_t has independent increments and continuous paths

Geometric Brownian motion

- $dS_t = \mu S_t dt + \sigma S_t dW_t$

Ito's Lemma

- Suppose X_t follows $dX_t = a(X_t, t)dt + b(X_t, t)dW_t$ and $f(X_t, t)$ is a twice-differentiable function. Then:

$$df = \left(\frac{\partial f}{\partial t} + a\frac{\partial f}{\partial x} + \frac{1}{2}b^2\frac{\partial^2 f}{\partial x^2}\right)dt + b\frac{\partial f}{\partial x}dW_t$$

Illustration of Brownian motion

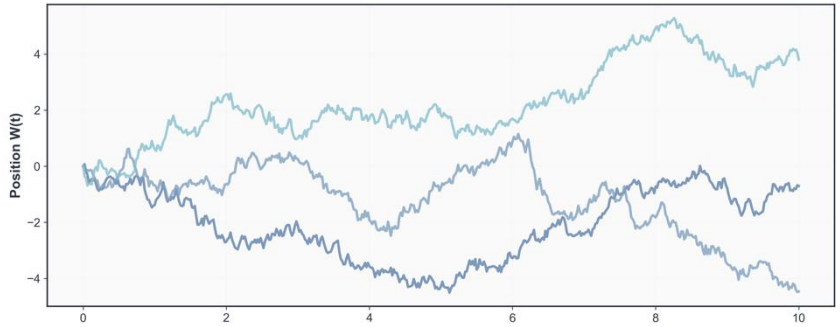
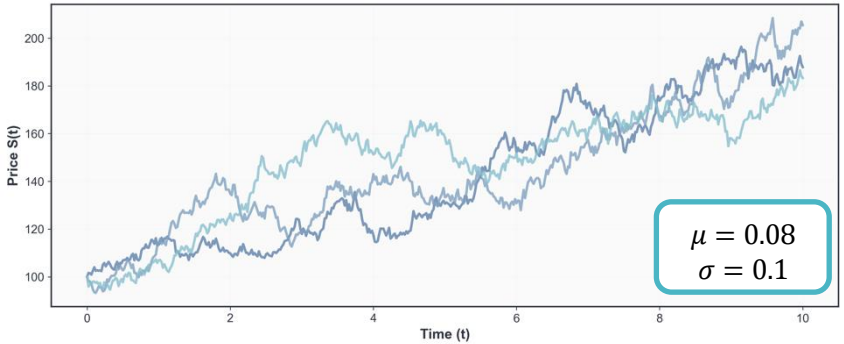


Illustration of Geometric Brownian motion



Black and Scholes model

For options pricing

The **Black-Scholes model**:

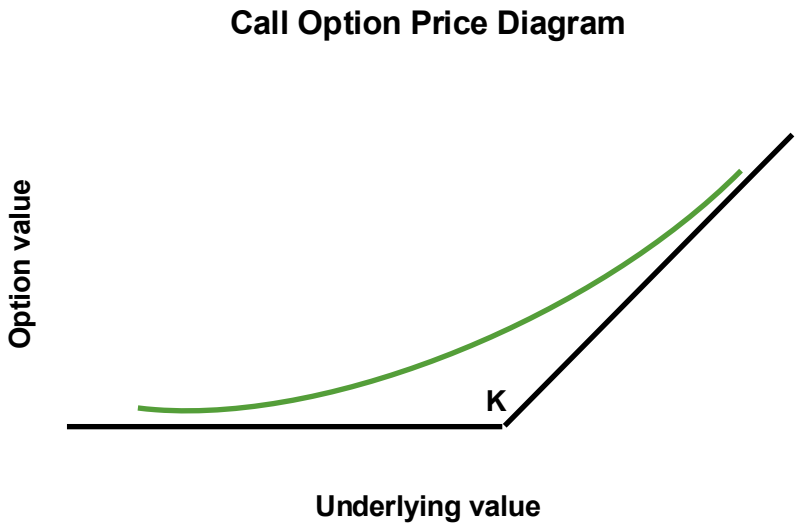
- is a **no-arbitrage framework** used to **price European-style options**
- assumes underlying is Brownian motion, trading in **continuous time** and have **no trading costs**
- is a **closed-form formula** to determine the fair value of a call or put option based on K, S, σ, r, T

$$C(S_t, t) = N(d_1)S_t - N(d_2)Ke^{-rt}$$

$$d_1 = \frac{1}{\sigma\sqrt{T-t}}\left[\ln\frac{S_t}{K} + \left(r + \frac{\sigma^2}{2}\right)(\mathrm{T} - \mathrm{t})\right]$$

$$d_2 = d_1 - \sigma\sqrt{T-t}$$

Option **Greeks** are the partial derivatives of the price with respect to the parameter values



Monte Carlo simulations

One of the most important tools in financial engineering

Monte Carlo simulation is a **numerical technique** used to **estimate expected values** by simulating many scenarios, it is used to **price derivatives** when closed-form solutions are not available

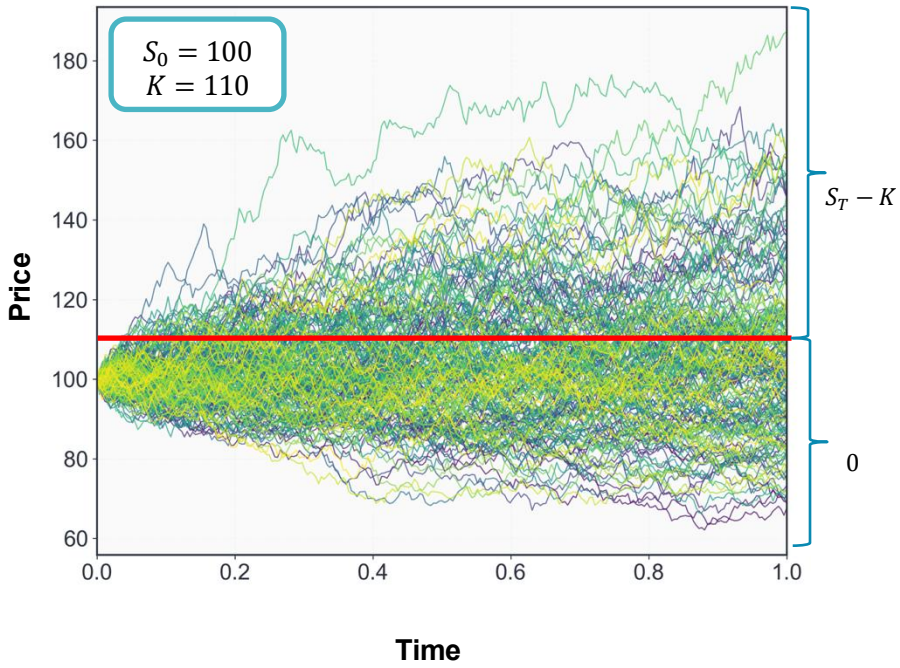
Applications

- Pricing of exotic **options**, credit **derivatives**, and structured products
- **Risk assessment**: VaR, Conditional VaR
- Stress testing and **scenario analysis**

Methodology

1. Calibrate stochastic process to the asset of interest
2. Simulate many price paths for the asset
3. Evaluate the payoff for each simulated path
4. Compute the average payoff and discount it to present value

Call option pricing example with MC simulations



Almgren and Chriss

An optimal execution algorithm

Almgren and Chriss is a model which defines the optimal execution strategy given the following assumptions on the price impact:

- $P_k = P_{k-1} + \theta v_{k-1} + \eta_{k-1}$ where $\eta \sim \mathrm{IID}(0, \sigma)$, where θv_{k-1} is a **linear market impact** term
- $\bar{P} = P_k + \rho v_k + sign(v_k)\frac{S}{2}$ where S is the bid ask spread and ρv_k is a **linear temporary market impact** term

The **execution costs** are defined as:

- $C(v) = \sum_{k=0}^{N-1} v_k \bar{P} - XP_0$

The objective function considers the trader's **risk aversion** λ and is defined as:

- $argmin_v\ \mathbb{E}[C(v)] - \lambda \mathrm{Var}(C(v))$

This problem has an analytical solution:

- $v_k = A\cosh(\sqrt{\frac{\lambda\sigma^2}{\rho}}(T - t_k))$

Illustration simulated impact

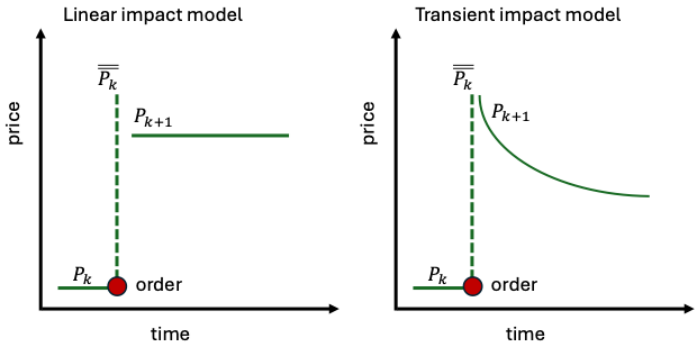
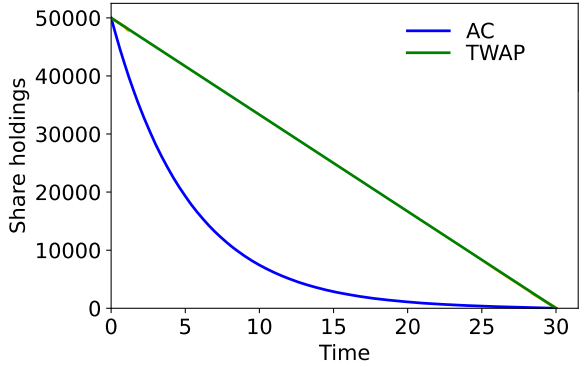


Illustration of Almgren Chriss execution trajectory



Avellaneda Stoikov model

Stochastic control framework for market making in a LOB

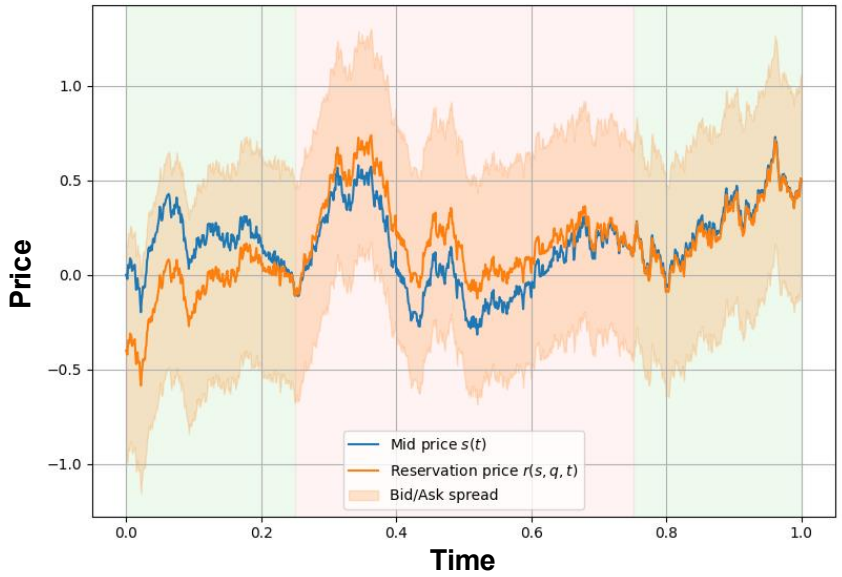
Assumptions:

- The mid-price follows a **Brownian motion**
- Market maker has **exponential utility function**: $U(W_T) = -e^{-\gamma W_T}$ where W_T is terminal wealth and γ is risk aversion
- Arrival rates of buy/sell market orders are **exponentially decreasing functions** of the distance between quoted prices and the mid-price: $Ae^{-k(q-m)}$

The market maker is continuously quoting a bid and ask prices calculated as a reservation price and plus the spread where:

- The reservation price is an adjustment to the mid-price, which accounts for the inventory held by the agent: $r = s - q\gamma\sigma^2(T-t)$
- The bid ask spread around the reservation price is: $\text{spread} = \gamma\sigma^2(T-t) + \frac{2}{\gamma}\ln(1+\frac{\gamma}{k})$

Simulation of reservation price and spread



Where we're going today

1	**Can we predict the future?**	Prediction vs trading
2	**Capital Markets landscape**	Who trades, why, and where you fit
3	**Financial engineering foundations**	Models for uncertainty, pricing & execution
4	**Systematic strategies**	From market intuition to backtested P&L
5	**Portfolio optimization**	Combining many imperfect strategies
6	**Machine learning in trading**	Extending the workflow — without the hype
7	**Gen AI on equity market neutral**	Using Gen AI on single stocks

Introduction to Systematic Trading

Defining and building a systematic trading strategy

Systematic trading definition

Quantitative trading uses mathematical and statistical models to identify trading opportunities

Common systematic trading strategies

- Trend *Focus next*
- Mean-reversion
- Seasonality
- Statistical arbitrage
- Market making
- Alternative data

Building a systematic trading strategy

A Objective	Financial assets, frequency, style
B Data	Price, LOB, sentiment, fundamental, economic
C Strategy	Define trading rules for the strategy
D Testing	Performance evaluation on historical data (backtesting)
Production	Connect to market via APIs, deploy strategy on a server

Trend strategy (1/3)

Are there any clearly visible patterns in the training set?

Strategy description

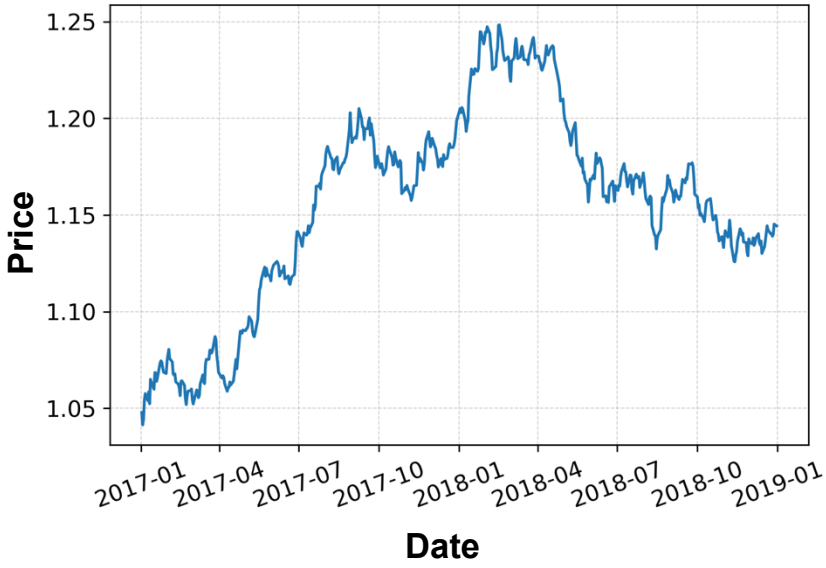
A Objective

- **Asset**: FX currency pair
- **Frequency**: daily
- **Style**: long/short

B Data

- **Data:** prices

Price process of asset on training set



Trend strategy (2/3)

Are there any clearly visible patterns in the training set?

Strategy description

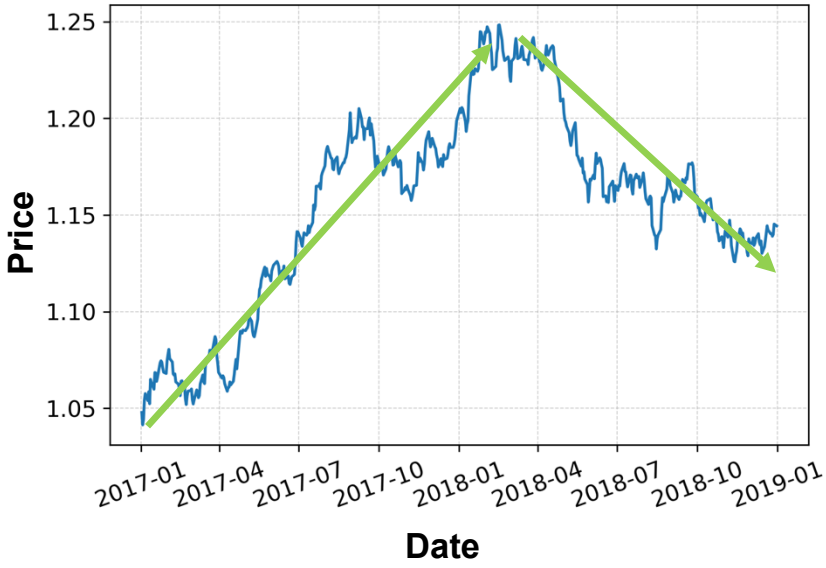
A Objective

- **Asset**: FX currency pair
- **Frequency**: daily
- **Style**: long/short

B Data

- **Data:** prices

Price process of asset on training set



Trend strategy (3/3)

Choose parameters on training set and analyze P&L on testing set

Strategy description

A Objective

- **Asset**: FX currency pair
- **Frequency**: daily
- **Style**: long/short

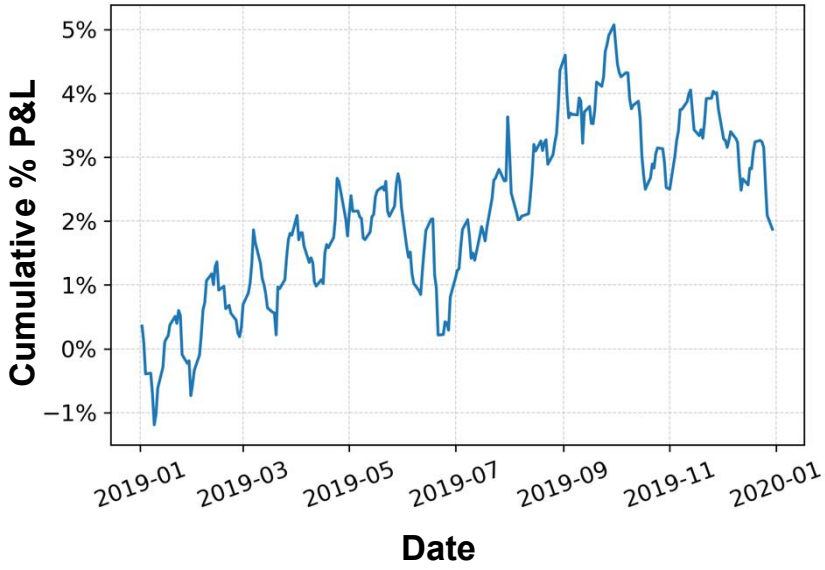
B Data

- **Data:** prices

C Strategy

+1 if $\mathrm{MA}_{\mathrm{short}} > \mathrm{MA}_{\mathrm{long}}$
-1 if $\mathrm{MA}_{\mathrm{short}} < \mathrm{MA}_{\mathrm{long}}$
where MA is the moving average

D Testing Cumulative % P&L on 2019



Judge a strategy in the right order

Build the picture step by step — the Sharpe ratio comes last, not first

1	**P&L**	The raw money made or lost over time
2	**Return**	P&L as a % of capital — so strategies of different sizes compare
3	**Volatility**	How much those returns bounce around: the risk
4	**Drawdown**	The worst peak-to-trough loss — what actually hurts

5 Sharpe ratio = return ÷ volatility

Return per unit of risk — a useful summary, but blind to tail risk and hidden dangers. A starting point, not proof.

Performance metrics

The basic performance metrics to evaluate a strategy

Performance metrics

- **Annualized P&L:**

$$\text{mean}\left(P\&L_{daily}\right) * 252$$

- **Annualized vol:**

$$\text{std}\left(P\&L_{daily}\right) * \sqrt{252}$$

- **Sharpe ratio:**

$$\frac{P\&L}{vol}$$

- **Maximum drawdown:**

$$Trough - Peak$$

Momentum strategy metrics on 2019

- **Annualized P&L:** 1.9%
- **Annualized vol:** 4.6%
- **Sharpe ratio:** 0.41
- **Maximum drawdown**: -3.3%

Mean reverting strategy (1/4)

Observe the average daily price movement with confidence intervals

Strategy description

A Objective

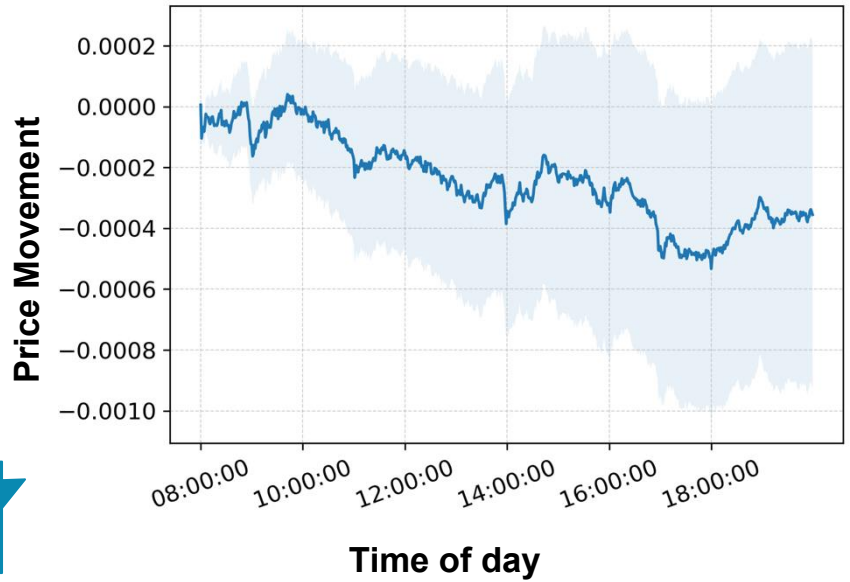
- **Asset**: FX currency pair
- **Frequency**: intraday
- **Style**: long/short

B Data

- **Data:** prices

The wide confidence intervals around the average asset movement suggest a mean reverting approach

Average asset movement



Mean reverting strategy (2/4)

The strategy achieves a surprisingly good result in the testing set

Strategy description

A Objective

- **Asset**: FX currency pair
- **Frequency**: intraday
- **Style**: long/short

B Data

- **Data:** prices

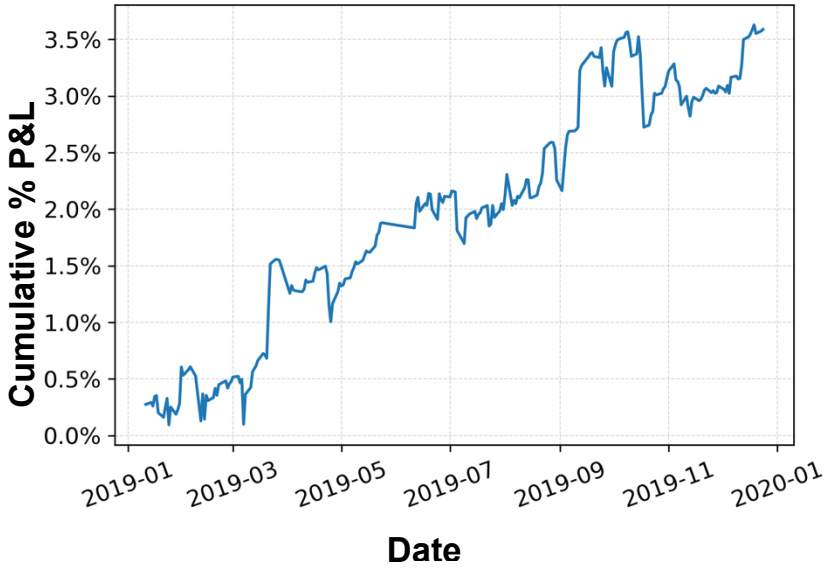
C Strategy

- Positioning: $-\sum_{i=0}^{T-2} f(i) R_{t-i}$

Legend

T = time horizon in minutes
R = returns

D Testing Cumulative % P&L on 2019



Performance Metrics

- Annualized return: 3.8%
- Sharpe: 2.03
- Max drawdown: -0.84%

Mean reverting strategy (3/4)

Each trades generates a cost proportional to the trade size

Defining Transaction Costs

- mid price $= \frac{1}{2}(\text{best ask} + \text{best bid})$
  - 4044.75
- spread $= (\text{best ask} - \text{best bid})$
  - 0.5
- transaction costs $= \text{trade size} * \frac{1}{2}\text{spread}$
- step P&L $= \text{position} * \text{market movement} - \text{transaction costs}$

Example of LOB

BID		ASK	
Volume	Price	Price	Volume
136	4044.50	4045.00	62
327	4044.00	4045.50	293
348	4043.50	4046.00	427
620	4043.00	4046.50	426
358	4042.50	4047.00	463
330	4042.00	4047.50	348
325	4041.50	4048.00	327
318	4041.00	4048.50	294
305	4040.50	4049.00	281
512	4040.00	4049.50	288

Mean reverting strategy (4/4)

Adding trading costs causes the performance to degrade

Strategy description

A Objective

- **Asset**: FX currency pair
- **Frequency**: intraday
- **Style**: long/short

B Data

- **Data:** prices

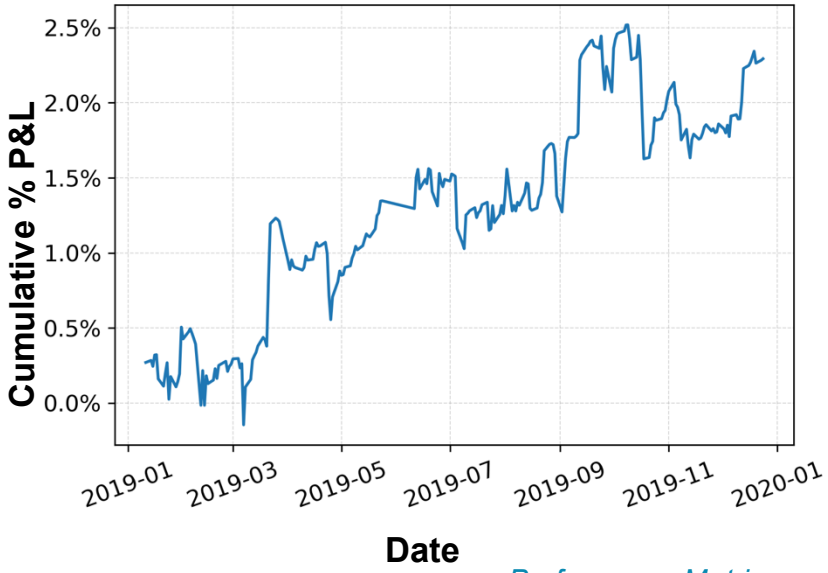
C Strategy

- Positioning: $-\sum_{i=0}^{T-2} f(i) R_{t-i}$

Legend

T = time horizon in minutes
R = returns

D Testing Cumulative % P&L on 2019

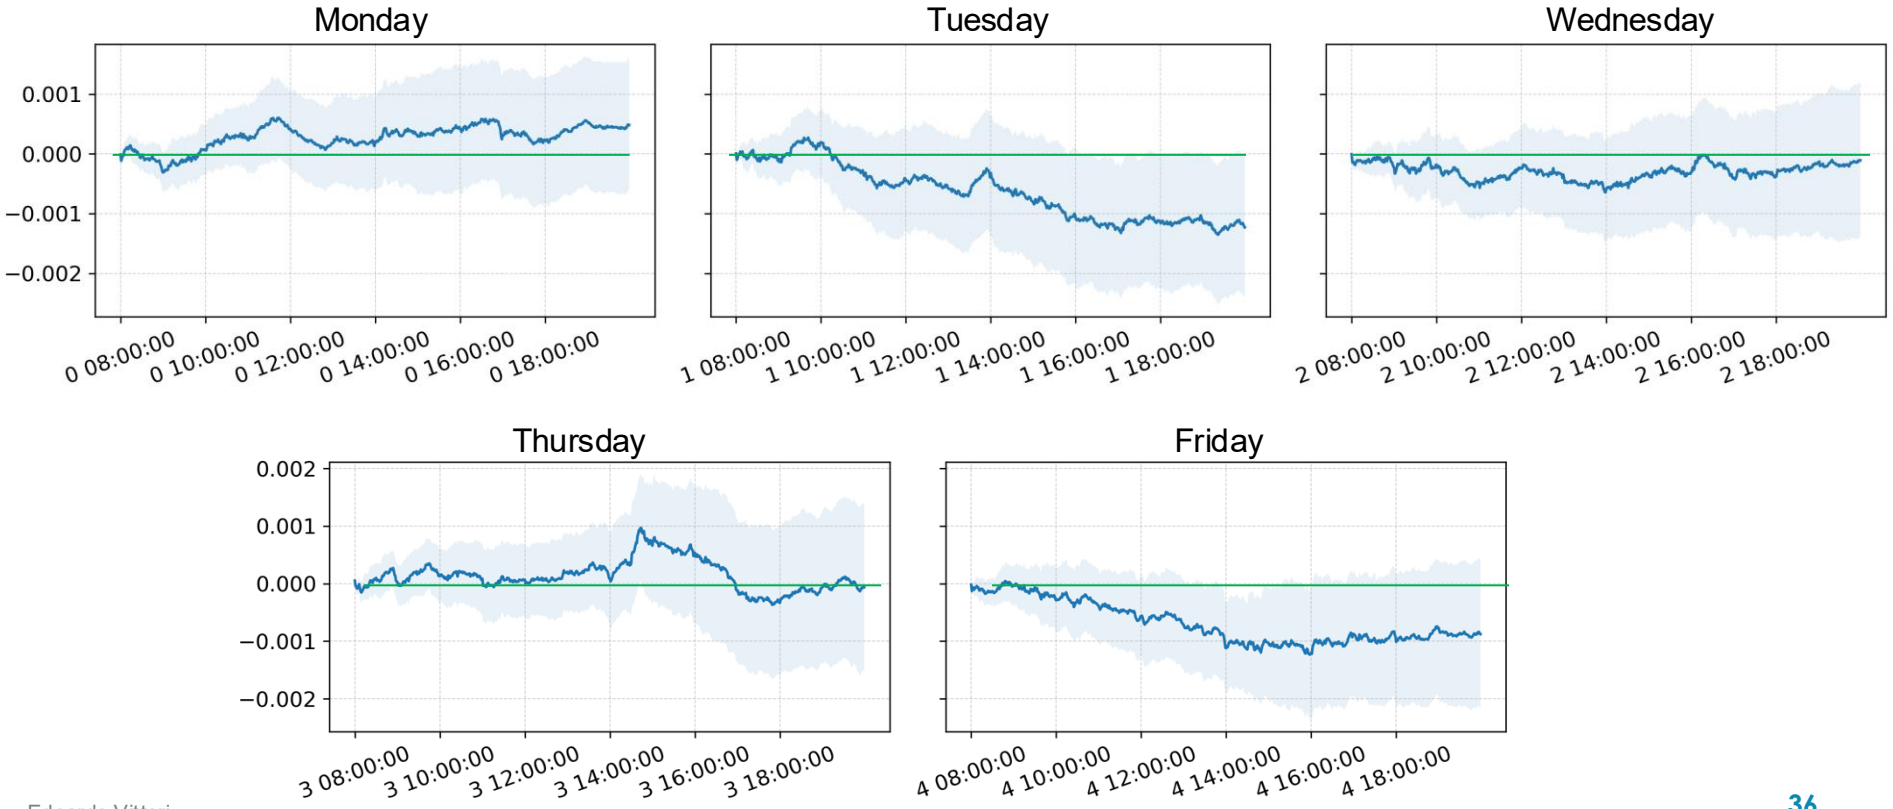


Performance Metrics

- Annualized return: 2.4%
- Sharpe: 1.2
- Max drawdown: -0.89%

Intraday seasonality strategy (1/2)

Analyze the average behavior on each day of the week



Intraday seasonality strategy (2/2)

Finding intervals where the cumulated intraday price is significantly different from zero

Strategy description

A Objective

- **Asset**: FX currency pair
- **Frequency**: intraday
- **Style**: long/short

B Data

- **Data:** prices

C Strategy

If on the training set the return is > 0 with statistical significance: long, elif it is < 0: go short

D Testing Cumulative % P&L on 2019

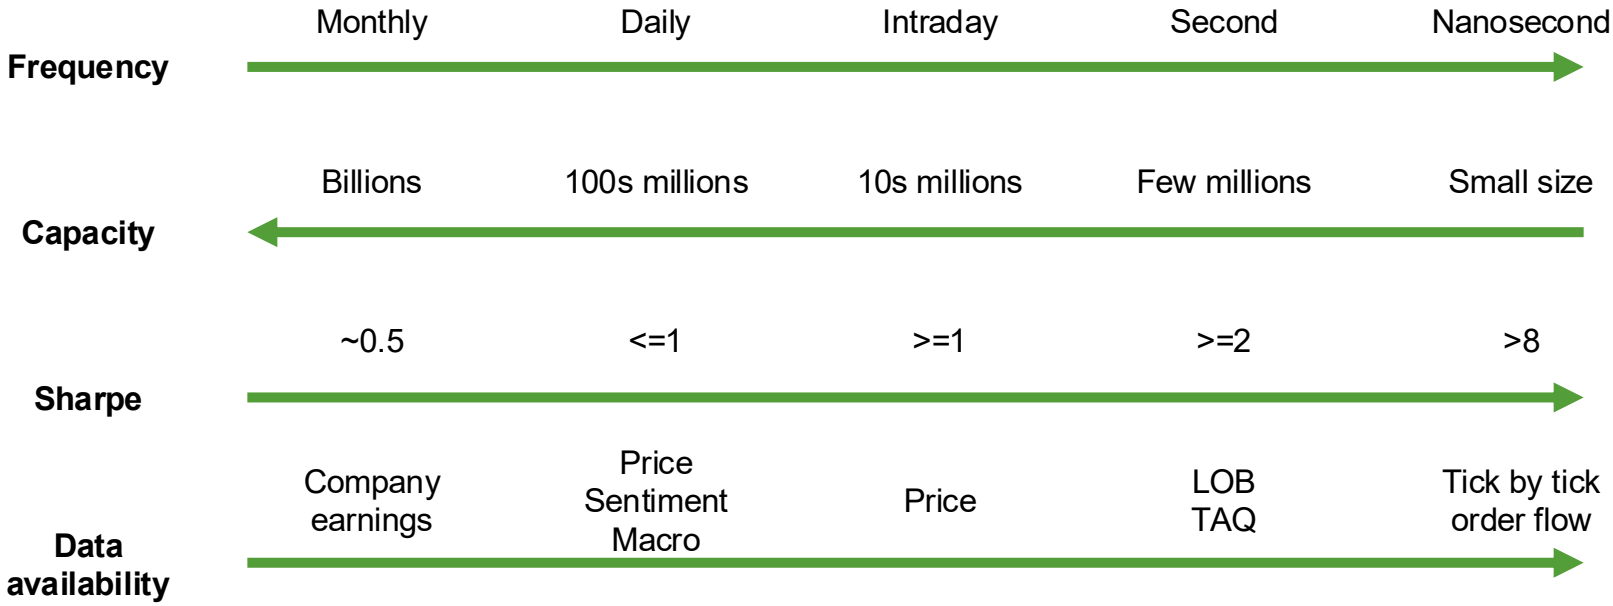


Performance Metrics

- Annualized return: 3.0%
- Sharpe: 1.06
- Max drawdown: -1.64%

Industry standards trading strategies

Higher frequencies mean higher infrastructure costs



Reinforcement Learning requires hundreds of thousands of datapoints to learn a trading strategy, which requires an intraday frequency. High frequency trading requires extremely expensive infrastructure. It is necessary to find a compromise

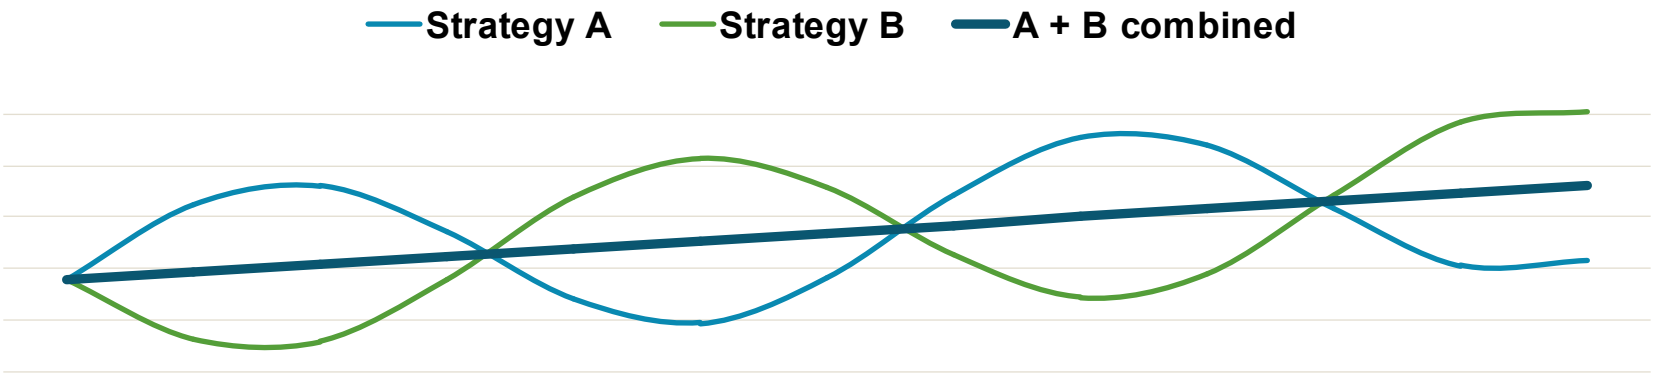
Where we're going today

1	**Can we predict the future?**	Prediction vs trading
2	**Capital Markets landscape**	Who trades, why, and where you fit
3	**Financial engineering foundations**	Models for uncertainty, pricing & execution
4	**Systematic strategies**	From market intuition to backtested P&L
5	**Portfolio optimization**	Combining many imperfect strategies
6	**Machine learning in trading**	Extending the workflow — without the hype
7	**Gen AI on equity market neutral**	Using Gen AI on single stocks

Diversify behavior

Two mediocre strategies that zig when the other zags can beat one great strategy alone

Diversification is not owning many things. It is owning return streams that do not move together — so a bad month for one is often an average month for the portfolio.



The lower the correlation between return streams, the more the combined volatility falls — for free. Correlation, not count, is what diversifies.

Combining strategies (1/2)

By combining uncorrelated strategies, it is possible to improve the Sharpe ratio

Sharpe of sum of two strategies

- $\text{Sharpe}(S_1) = \frac{\mu_1}{\sigma_1}$, $\text{Sharpe}(S_2) = \frac{\mu_2}{\sigma_2}$
- $\text{Sharpe}(S_1 + S_2) = \frac{\mu_1+\mu_2}{\sqrt{\sigma_1^2+\sigma_2^2+2\rho\sigma_1\sigma_2}}$

Correlation: ρ	Sharpe(S1+S2)
1	$\frac{\mu_1+\mu_2}{\sigma_1+\sigma_2}$
0	$\frac{\mu_1+\mu_2}{\sqrt{\sigma_1^2+\sigma_2^2}}$
-1	$\frac{\mu_1+\mu_2}{\lvert\sigma_1-\sigma_2\rvert}$

Since $(\sigma_1+\sigma_2)^2 > \sigma_1^2+\sigma_2^2$, uncorrelated strategies are preferrable

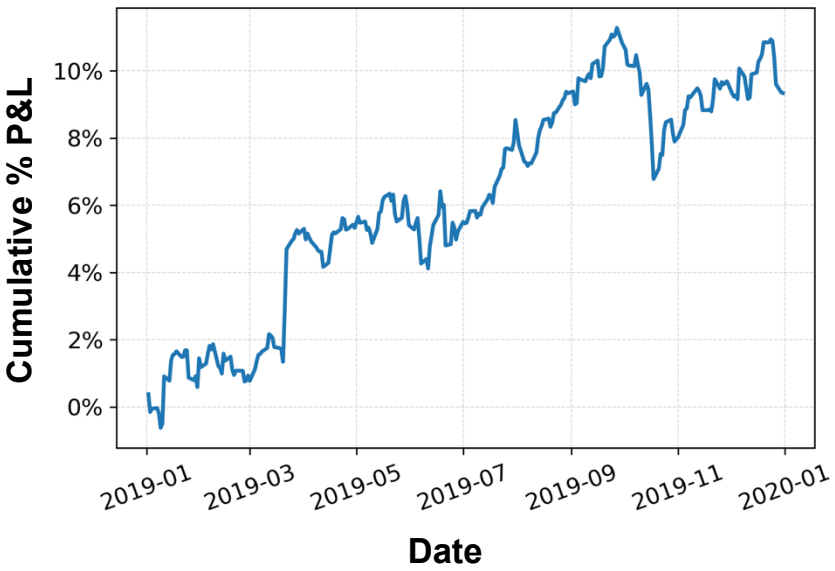
Correlation of our strategies

	Momentum	Mean reverting	Seasonality
Momentum	1.00	-0.18	0.25
Mean reverting	-0.18	1.00	-0.10
Seasonality	0.25	-0.10	1.00

Combining strategies (2/2)

A naive approach is to give the same weight to each strategy

Cumulative % P&L on 2019



Performance metrics

	Sharpe	Vol
Aggregated	1.6	0.06
Momentum	0.4	0.10
Variance	1.2	0.10
Seasonality	1.1	0.10

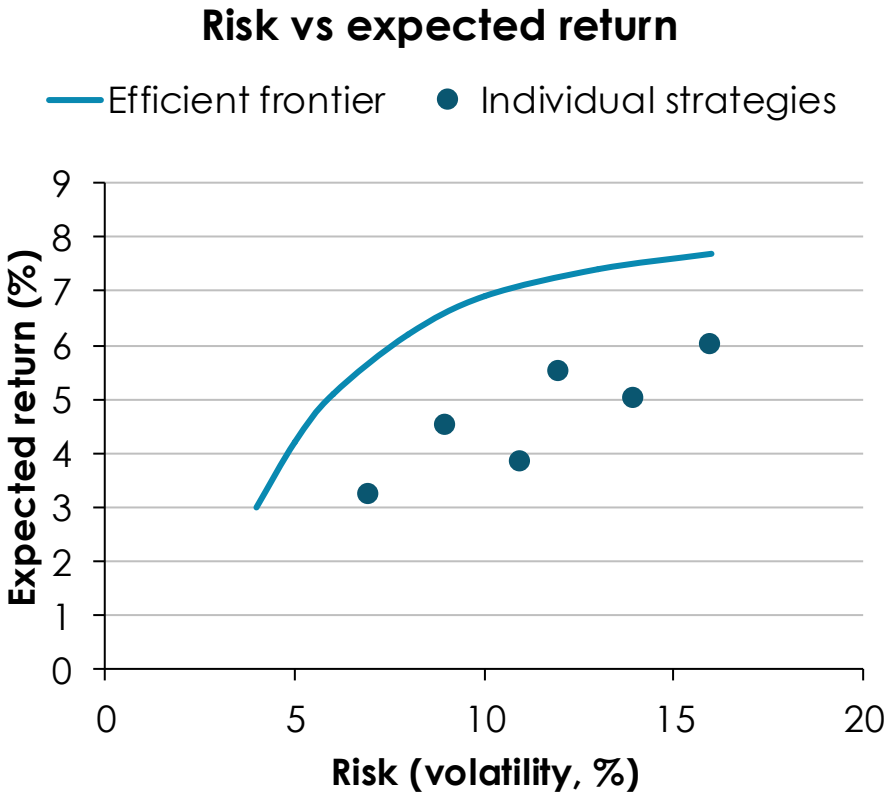
We can see that by averaging the strategies, the Sharpe improves as volatility has decreased. Since we are comfortable with a 0.10 vol, we can increase the size of the strategy.

The efficient frontier

Mean-variance optimization: the best expected return available at each level of risk

For a target risk, which mix of strategies gives the highest expected return?

- Inputs: expected returns, volatilities and the correlation matrix
- Output: the portfolio weights
- The frontier is the set of best possible trade-offs
- There is no single best portfolio — it depends on how much risk you can hold



Portfolio optimization in practice

Setting weights by risk, not by return forecasts

Set weights by risk, not by forecasting returns. Each method leans on the covariance matrix to a different degree — trading simplicity for precision.

Four risk-based approaches, simplest to most complex:

Inverse volatility

Higher volatility, smaller weight. Ignores correlations — simple and robust.

Minimum variance

Lowest-volatility weights; leans hard on the covariance matrix.

Equal risk contribution

Every strategy contributes equal risk.

Risk budgeting

Give each strategy a target share of total risk; solve for those weights.

Where we're going today

1	**Can we predict the future?**	Prediction vs trading
2	**Capital Markets landscape**	Who trades, why, and where you fit
3	**Financial engineering foundations**	Models for uncertainty, pricing & execution
4	**Systematic strategies**	From market intuition to backtested P&L
5	**Portfolio optimization**	Combining many imperfect strategies
6	**Machine learning in trading**	Extending the workflow — without the hype
7	**Gen AI on equity market neutral**	Using Gen AI on single stocks

Why financial ML is genuinely hard

The tools that master images and language struggle on markets

Low signal-to-noise

Most price moves are noise, not signal

Non-stationarity

The rules change over time

Limited data

Decades of history, not billions of samples

Competition

Others erase edges as you find them

Costs & implementation

A paper edge can die after costs

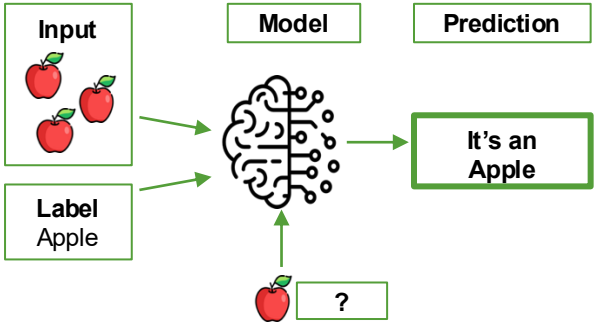
Regime change

What worked can suddenly stop working

Introduction to Machine Learning

There are three main paradigms in ML

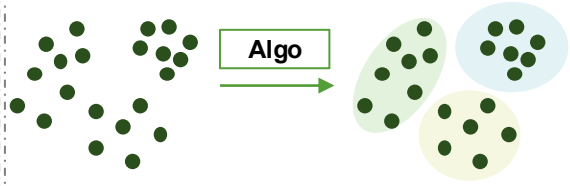
Supervised Learning



Usage

- Regression
- Classification

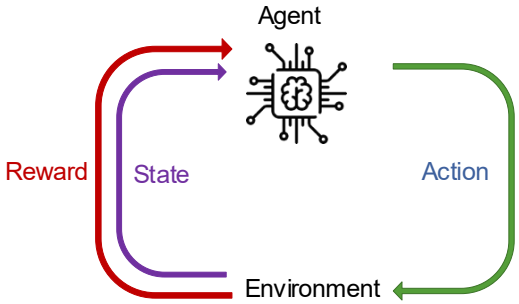
Unsupervised Learning



Usage

- Clustering
- Dimensionality reduction
- Feature extraction

Reinforcement Learning

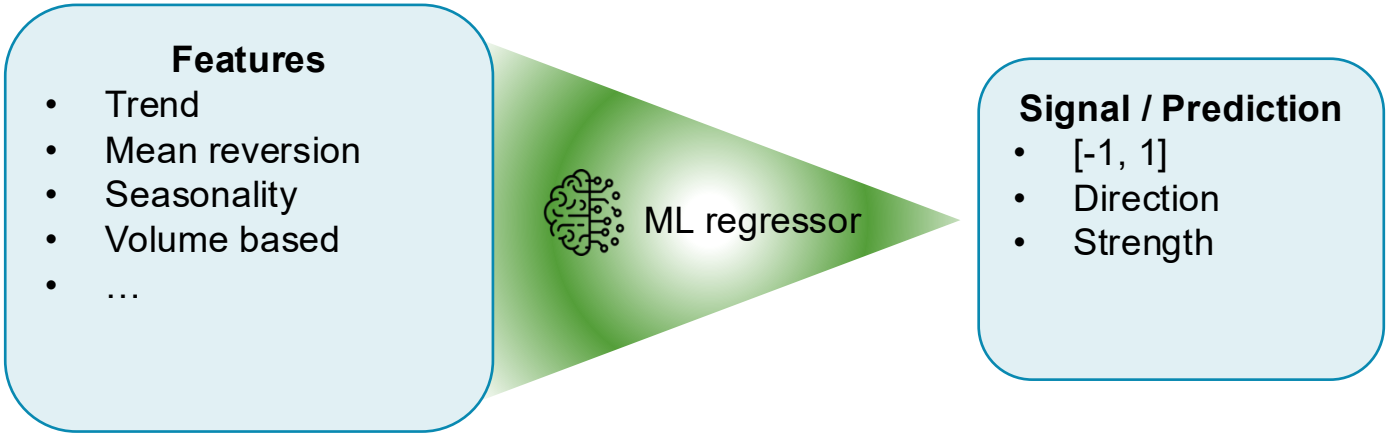


Usage

- Optimal Control
- Game playing

Supervised learning in trading

Learn the map from market features to a signal



Model choices

Linear / Ridge · Gradient-boosted trees · Random forest · Small neural nets — simpler models often win on noisy, low-signal financial data

From data to generalization

Avoiding overfitting in ML

Dataset

A dataset is a collection of **examples** with:

- **Input data (features)**: image, text, numbers
- **Labels (targets)**: the correct answer (class or value)

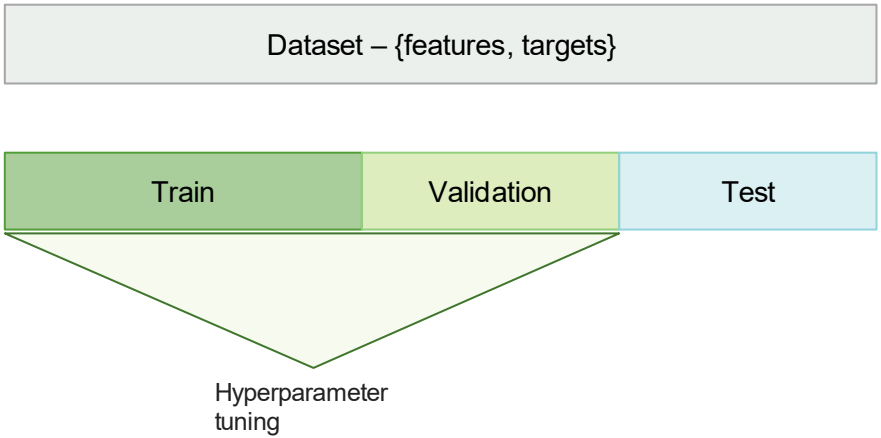
Data Processing

Normalize, clean, encode, or scale input features, then split the data into training set, validation set, testing set.

Key concepts

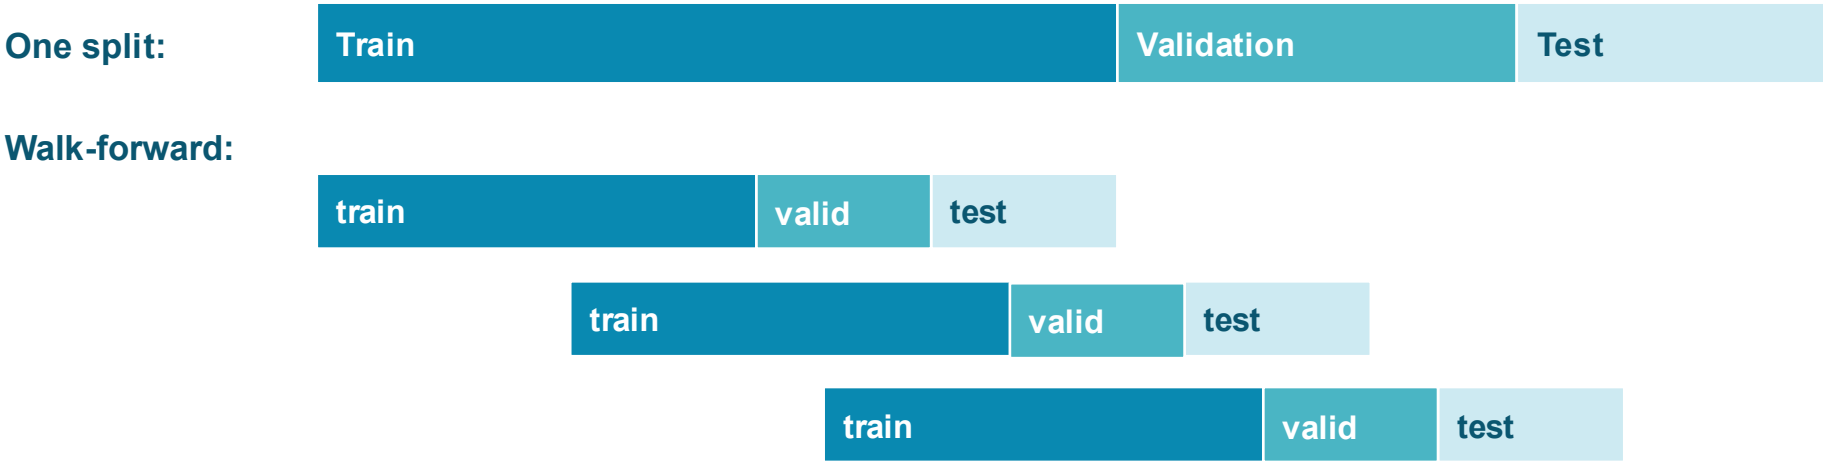
- **Overfitting**: When the model learns **too much from the training data**, and performs poorly on test data
- **Regularization** are techniques to **reduce overfitting** by simplifying the model such as adding an **L2 / L1 penalty, using dropout or early stopping**
- **Hyperparameter tuning** is the process of selecting the best set of hyperparameters for a ML model

Dataset splitting illustration



Validate by walking forward in time

Random shuffling leaks the future — always test on data the model has never seen



Time → each window trains on the past, tests on the next unseen slice, then rolls forward.

Include realistic costs from the very first backtest — not bolted on at the end.

Hyperparameter Tuning

Training these models make take a large amount of time and compute power



References

- Akiba, Takuya, et al. "Optuna: A next-generation hyperparameter optimization framework." *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 2019
- Bergstra, James, Dan Yamins, and David D. Cox. "Hyperopt: A Python Library for Optimizing the Hyperparameters of Machine Learning Algorithms." *SciPy* 13 (2013): 20.
- Hutter, Frank, Holger H. Hoos, and Kevin Leyton-Brown. "SMBO: Sequential model-based optimization for general algorithm configuration." *Learning and Intelligent Optimization: 5th International Conference*, 2011.

Reinforcement Learning for Systematic Trading

Problem description and MDP definition

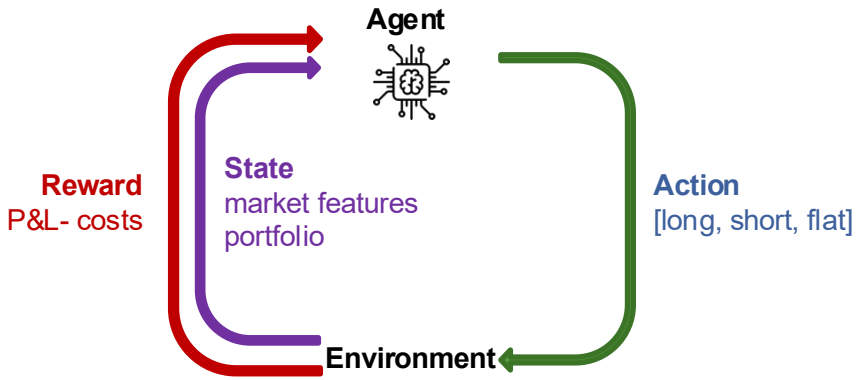
Definition

- At each timestep, decide whether to go long, short or flat to maximize gains

Markov Decision Process (MDP)

- **State:** market features, bid-ask spread, current portfolio, date/time
- **Action:** long, short, flat
- **Reward:** P&L – transaction costs

MDP Graphic



- The objective is finding the policy π which maximizes the discounted sum of the rewards
- $J_\pi = \mathbb{E}_\pi[\sum \gamma^t R_t]$

Q-function and Policy Search

RL algorithms enable the learning of the policy π

Q-learning

- Q-function

  $$Q_\pi = \mathbb{E}_\pi[\sum \gamma^t R_t | s_0, a_0]$$

- Bellman Equation

  $$Q_\pi = r(s,a) + \gamma \mathbb{E}_{s',a'}[Q_\pi(s',a')]$$

- Q-learning algorithm

  $$Q_t(s,a) = r(s,a) + \gamma \max_{a'} Q_t(s',a')$$

- Q-learning is a tabular algorithm which can be generalized using function approximators such as XGBoost

Algorithm examples

- DQN
- DDQN
- FQI

Policy Search

- Policy update

  $$\theta_{t+1} = \theta_t + \alpha \nabla_\theta J_{\pi_\theta}$$

- Policy gradient theorem

  $$\nabla_\theta J_{\pi_\theta} = \mathbb{E}[\nabla \log \pi_\theta(a|s) Q_{\pi_\theta}(s,a)]$$

- The policy is a parametric and differentiable function, usually a neural network

Algorithm examples

- REINFORCE
- TRPO
- PPO

Reinforcement Learning for FX trading (1/2)

Experimental results - performance

Strategy description

A Objective

- **Asset**: FX currency pair
- **Frequency**: intraday
- **Style**: long/short

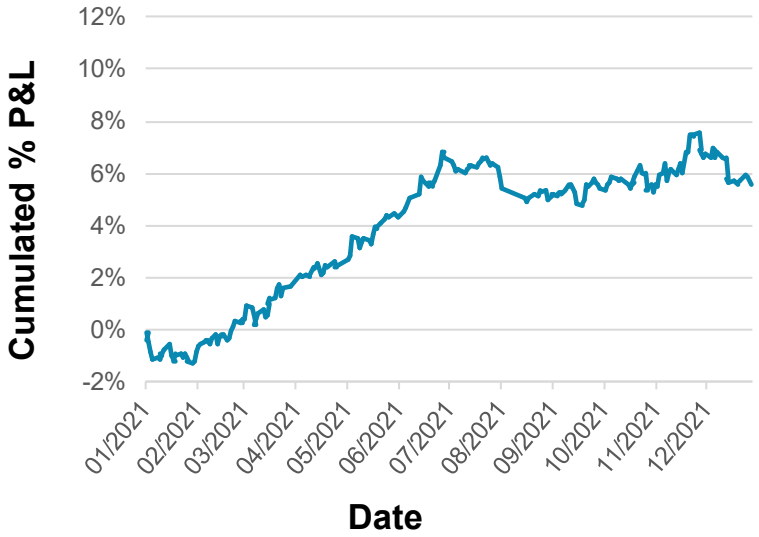
B Data

- **Data:** prices

C Strategy

Training with reinforcement learning on 2018-2019, validation on 2020

D Testing Cumulative % P&L on 2021



Performance Metrics

- Annualized return: 6.0%
- Sharpe: 1.7
- Max drawdown: -2.0%

Riva, Antonio, et al. "Learning FX trading strategies with FQI and persistent actions." *ICAIF* 2021.

Where we're going today

1	**Can we predict the future?**	Prediction vs trading
2	**Capital Markets landscape**	Who trades, why, and where you fit
3	**Financial engineering foundations**	Models for uncertainty, pricing & execution
4	**Systematic strategies**	From market intuition to backtested P&L
5	**Portfolio optimization**	Combining many imperfect strategies
6	**Machine learning in trading**	Extending the workflow — without the hype
7	**Gen AI on equity market neutral**	Using Gen AI on single stocks

Cross-sectional momentum strategy

Make money from relative performance between stocks

Strategy description

A Objective

- **Asset**: Single stocks, S&P500 universe
- **Frequency**: monthly
- **Style**: cross sectional, market neutral

B Data

- **Data**: stock prices

C Strategy

- At the end of each month, you calculate **each stock's total return** over the past 6 or 12 months

General rule for equity long short

You **rank all stocks in the universe from best to worst performers**

- Go **long** the 10% of top-ranked stocks
- Go **short** the 10% of bottom-ranked stocks
- Construct **dollar neutral** or **beta neutral** portfolio

Equity long-short market neutral strategies try to make money from relative performance between stocks, while keeping broad market exposure close to zero.

Value strategy

Buy cheap stocks and short expensive stocks

Strategy description

A Objective

- **Asset**: Single stocks, S&P500 universe
- **Frequency**: monthly
- **Style**: cross sectional, market neutral

B Data

- **Data:** fundamental data

C Strategy

Consider metrics such as:

- Book-to-price
- Earnings /price

General rule for equity long short

You **rank all stocks in the universe from best to worst performers**

- Go **long** the 10% of top-ranked stocks
- Go **short** the 10% of bottom-ranked stocks
- Construct **dollar neutral** or **beta neutral** portfolio

Equity long-short market neutral strategies try to make money from relative performance between stocks, while keeping broad market exposure close to zero.

Quality strategy

buying financially strong, stable companies and shorting weak, fragile companies

Strategy description

A Objective

- **Asset**: Single stocks, S&P500 universe
- **Frequency**: monthly
- **Style**: cross sectional, market neutral

B Data

- **Data:** fundamental data

C Strategy

Consider metrics such as:

- Profits / assets
- Earnings stability

General rule for equity long short

You **rank all stocks in the universe from best to worst performers**

- Go **long** the 10% of top-ranked stocks
- Go **short** the 10% of bottom-ranked stocks
- Construct **dollar neutral** or **beta neutral** portfolio

Equity long-short market neutral strategies try to make money from relative performance between stocks, while keeping broad market exposure close to zero.

Generative AI strategy

Can we create a better ranking using LLMs?

Strategy description

A Objective

- **Asset**: Single stocks, S&P500 universe
- **Frequency**: monthly
- **Style**: cross sectional, market neutral

B Data

- **Data:** prices, news, macro releases

C Strategy

Use **LLMs** to enhance a cross-sectional market neutral strategy by **including news, articles and macro data**

General rule for equity long short

You **rank all stocks in the universe from best to worst performers**

- Go **long** the 10% of top-ranked stocks
- Go **short** the 10% of bottom-ranked stocks
- Construct **dollar neutral** or **beta neutral** portfolio

Equity long-short market neutral strategies try to make money from relative performance between stocks, while keeping broad market exposure close to zero.

Generative AI in trading

Powerful for parsing unstructured data at scale — dangerous as “what should I buy?”



GOOD USES

Summarizing filings · classifying news · earnings-call tone · central-bank speeches · coding & workflow help

THE DANGER

Asking an LLM directly what to buy. Its output only counts once it is a measurable, tested, cost-aware signal — just like any other.

Leveraging agents for equity trading

System design, data requirements, and implementation considerations

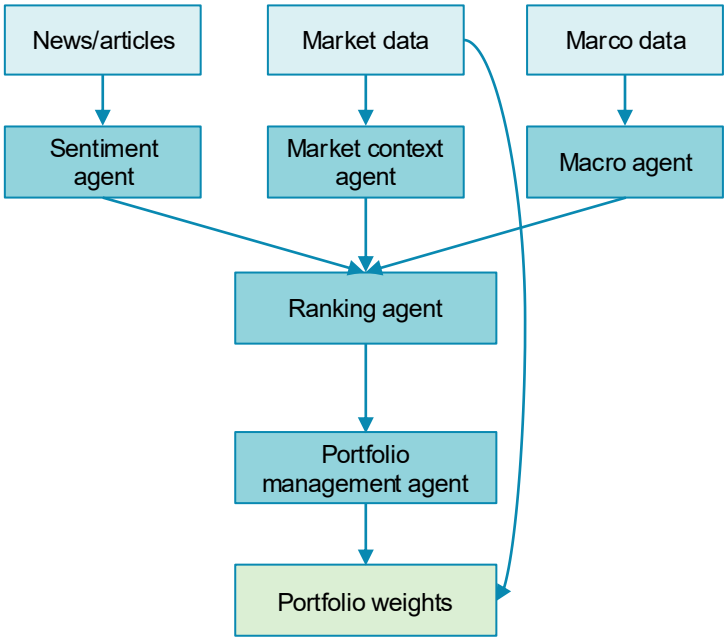
System design choices

- Choice of LLM and context length limitations
- Prompt engineering for financial understanding
- Fine-tuning vs. zero-shot vs. few-shot use
- Avoiding forward-looking bias (LLMs trained beyond back-test window)

Dataset

- **News/articles:** Company news, financial blogs, press releases
- **Market data:** Price history, volume, volatility, technical indicators
- **Macro data:** Interest rates, inflation, GDP, employment stats

Diagram of agent configuration



Five lessons to take away

1. Trading is decision-making under uncertainty.
2. A systematic trading idea must be a testable rule.
3. Backtests are useful but avoid overfitting.
4. Portfolios matter more than individual signals.
5. Do not use LLMs for ‘what should I buy?’

Quantitative Finance

Q&A

Edoardo Vittori
edov@me.com



The opinions expressed in this document are solely those of the authors and do not represent in any way those of their present and past employers