Quantitative Trading with Machine Learning

Edoardo Vittori
Vice President, CVA Management & AI Investments
IMI Corporate & Investment Banking Division, Intesa Sanpaolo

European Conference on Artificial Intelligence
Workshop on AI in Finance
October 19th 2024

Today's Agenda

Quantitative Trading

- Introduction
- Daily momentum
- Intraday mean reverting
- Intraday seasonality
- Strategy combination

RL for Quant Trading

- Reinforcement Learning for quantitative trading

Trading Futures with ML

- The Limit Order Book
- Feature engineering
- Strategy generation
- Optimizing execution

Today’s Agenda

Quantitative Trading

- Introduction
- Daily momentum
- Intraday mean reverting
- Intraday seasonality
- Strategy combination

RL for Quant Trading

- Reinforcement Learning for quantitative trading

Trading Futures with ML

- The Limit Order Book
- Feature engineering
- Strategy generation
- Optimizing execution

Introduction to quantitative trading

Defining and building a quantitative trading strategy

Quantitative trading definition

Quantitative trading uses mathematical and statistical models to identify trading opportunities

Common quantitative trading strategies

Focus next

- Momentum
- Mean-reversion
- Seasonality
- Statistical arbitrage
- Market making
- Alternative data

Building a quant 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
Production	Connect to market via APIs, deploy strategy on a server

Momentum strategy (1/2)

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

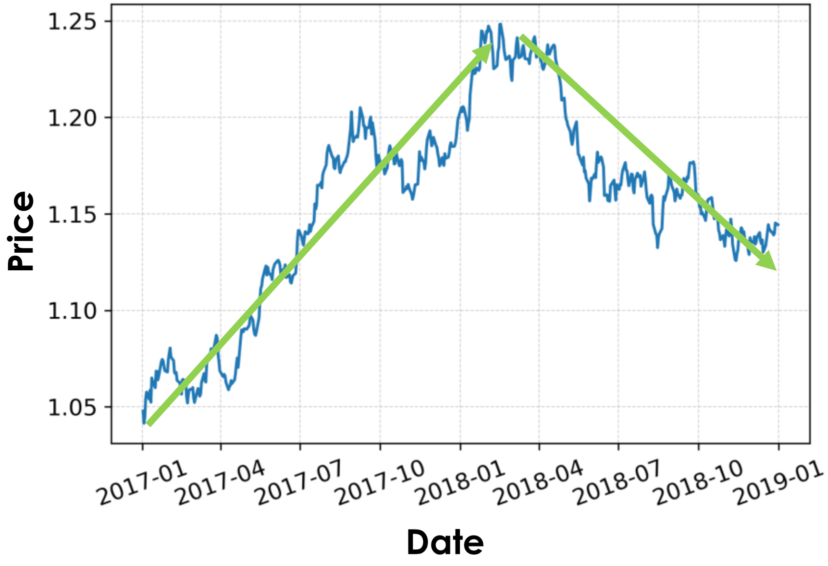

Momentum strategy (2/2)

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 $MA_{short} > MA_{long}$
-1 if $MA_{short} < MA_{long}$
where MA is the moving average

D Testing Cumulative % P&L on 2019

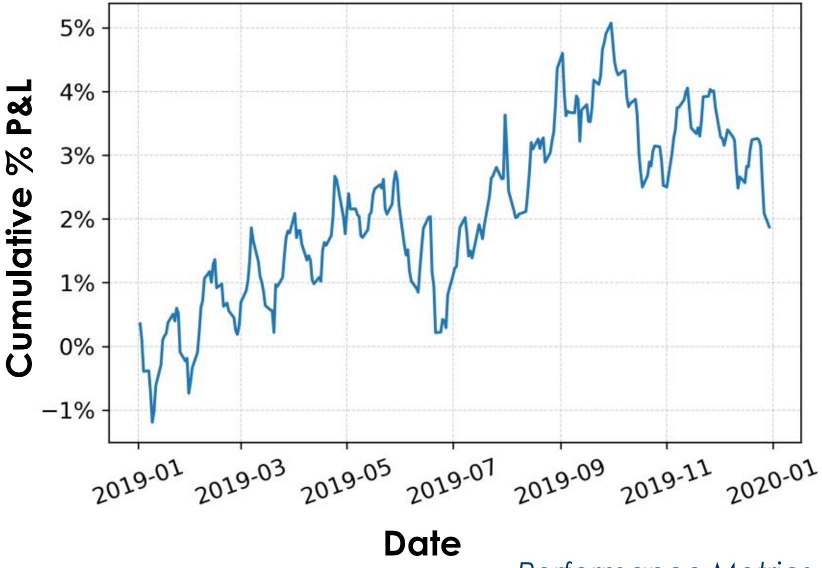

Performance Metrics

- Annualized return: 1.9%
- Sharpe: 0.41
- Max drawdown: -3.3%

Mean reverting strategy (1/3)

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

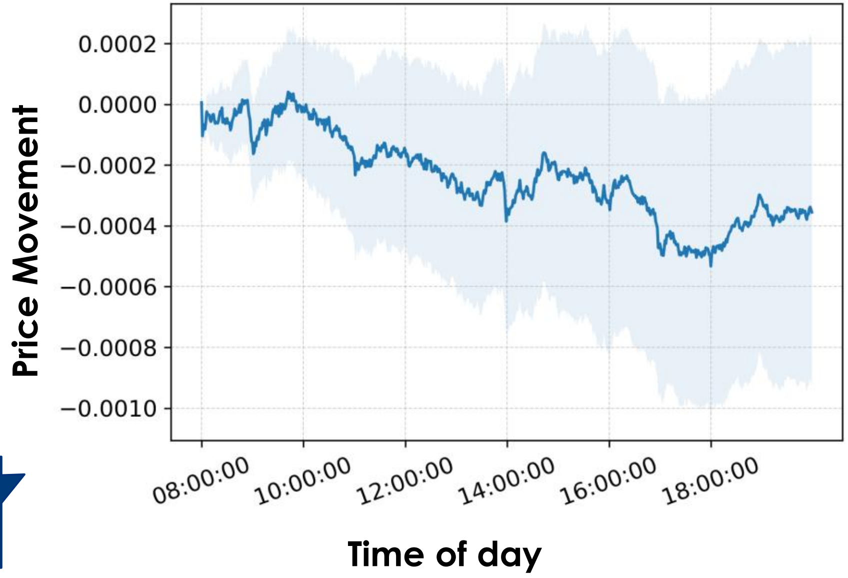

Mean reverting strategy (2/3)

The positive performance confirms a mean reverting behavior of the asset's price

Strategy description

A Objective

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

B Data

- **Data:** prices

C Strategy

Positioning: $-\sum_{i=0}^{T-2}(T-i-1)R_{t-i}$

$\mathbb{E}(P\&L) = \frac{1}{2}(T\mathrm{Var}_1 - \mathrm{Var}_\mathrm{T} - \mu^2 T(T-1))$

Legend

T = time horizon in minutes
R = returns
μ = average asset return
Var_1 = 1-period variance
Var_T = T-period variance

D Testing Cumulative % P&L on 2019

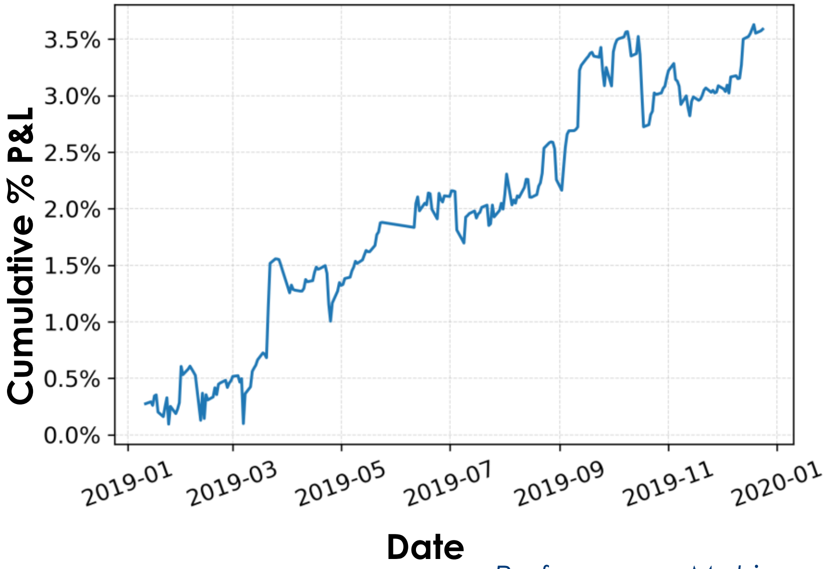

Performance Metrics

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

Mean reverting strategy (3/3)

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: $-\sum_{i=0}^{T-2}(T-i-1)R_{t-i}$

$\mathbb{E}(P\&L) = \frac{1}{2}(T\mathrm{Var}_1 - \mathrm{Var}_\mathrm{T} - \mu^2 T(T-1))$

Legend

T = time horizon in minutes
R = returns
μ = average asset return
Var_1 = 1-period variance
Var_T = T-period variance

D Testing Cumulative % P&L on 2019

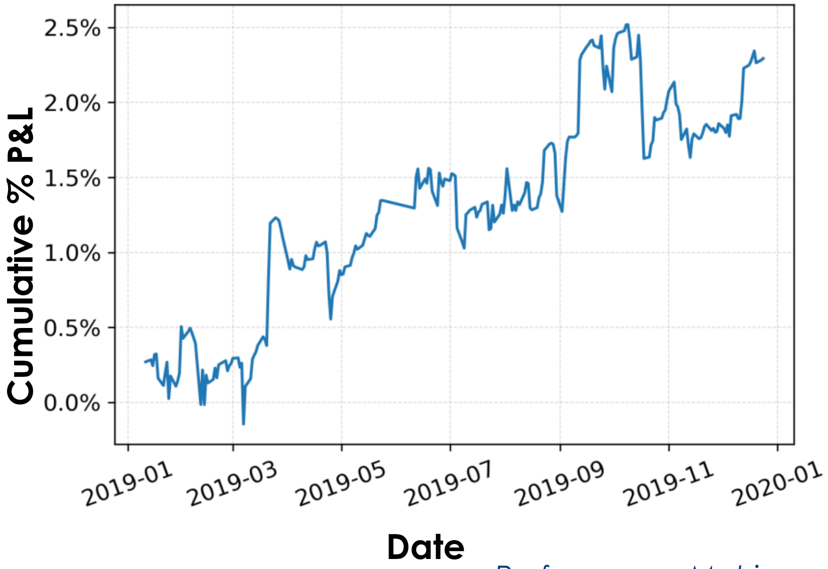

Performance Metrics

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

Intraday seasonality strategy (1/2)

Analyze average behavior on each day of the week

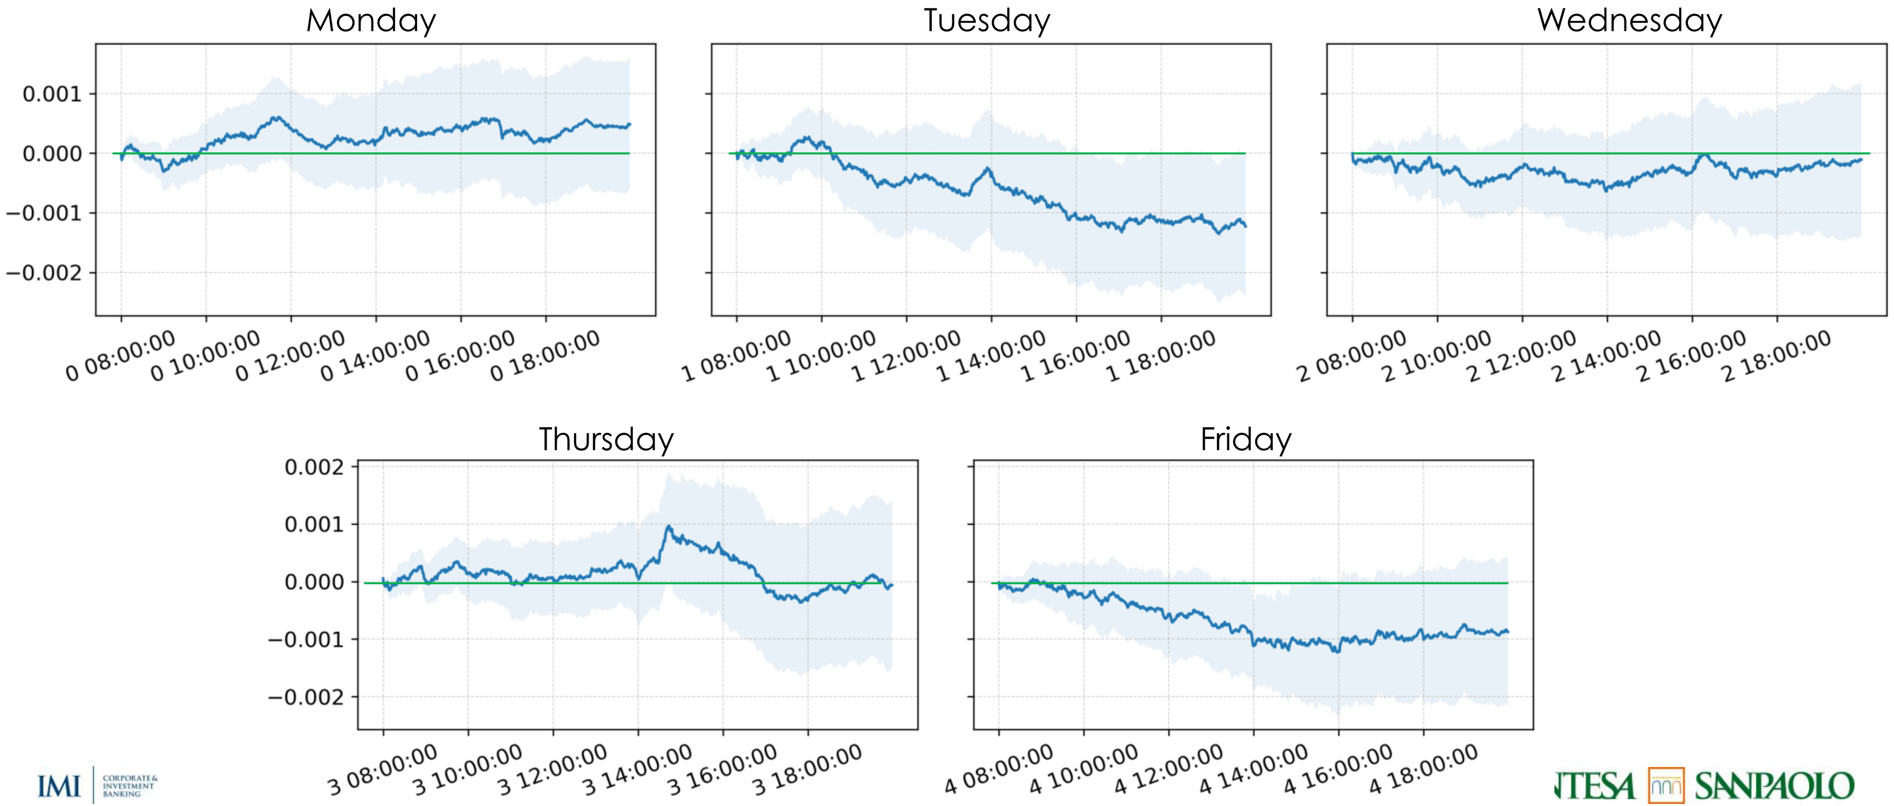

Intraday seasonality strategy (2/2)

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 price is > 0 with statistical significance: long, elif it is < 0: go short

D Testing Cumulative % P&L on 2019

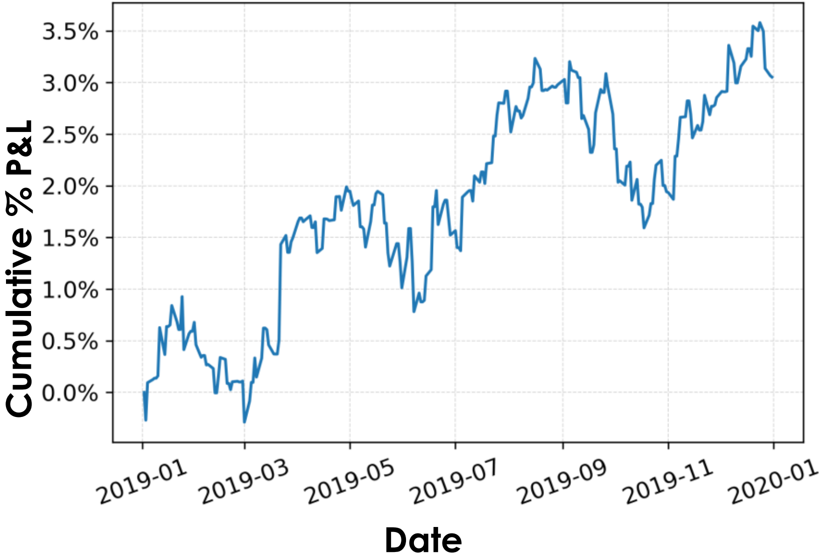

Performance Metrics

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

Combining strategies (1/3)

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

Combining strategies (2/3)

By changing the size of the strategies, we can modify the relative volatility

Strategies with same size

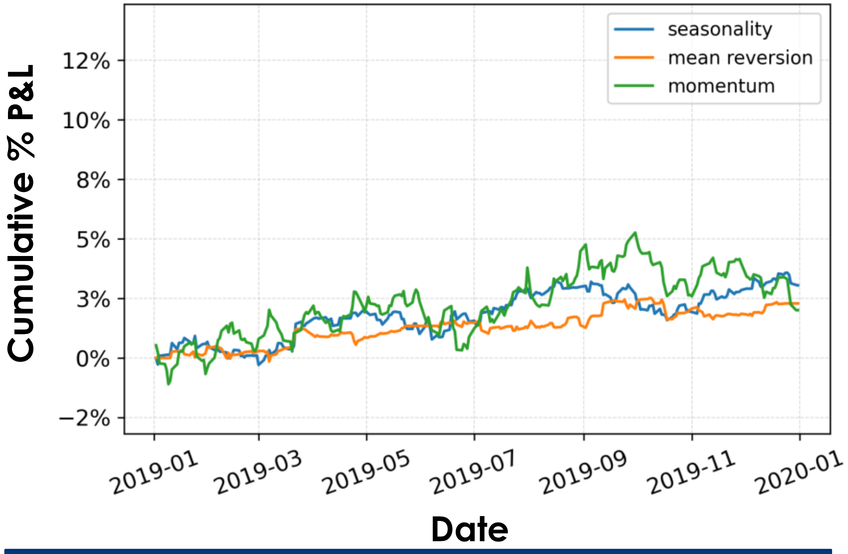

Strategies with equal vol

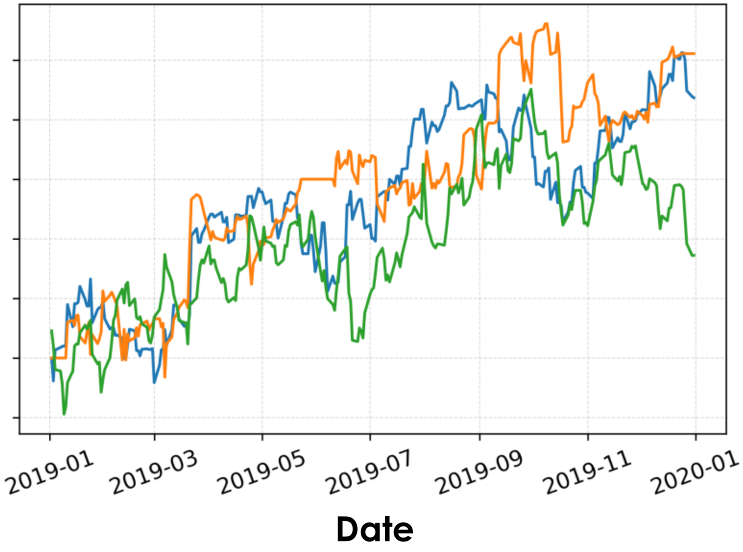

The size multiplier is $\boldsymbol{m_i} = \frac{\sigma}{\sigma_i}$ where σ_i is the volatility of strategy I and σ the target volatility. Choose a level of volatility you are comfortable with, 0.10 in the example

	Momentum	Mean reversion	Seasonality
m_i	2.1	5.6	3.6

Combining strategies (3/3)

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

Cumulative % P&L on 2019

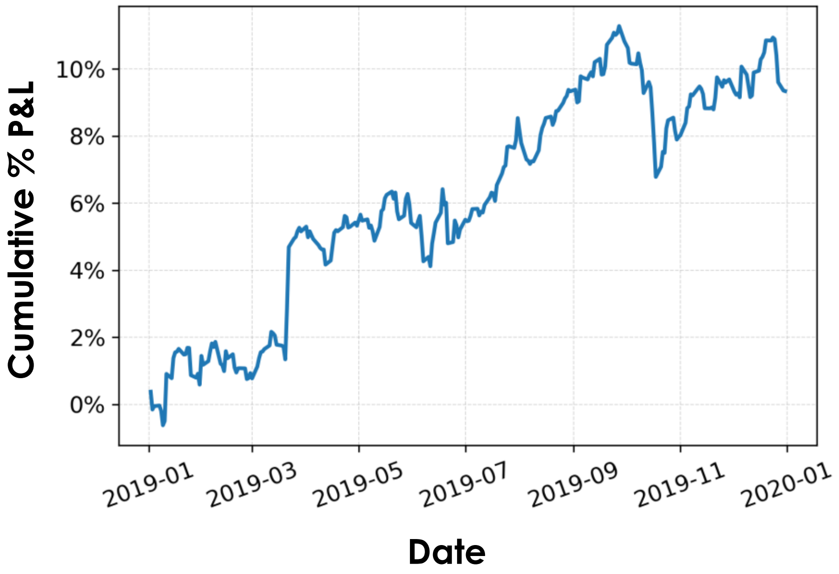

Performance metrics

	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

Expert Learning (1/3)

A dynamic, data driven approach to portfolio optimization

Characteristics

- While classical portfolio optimization models usually necessitate an estimation of the covariance matrix and the returns, **expert learning algorithms** are **data driven**
- Field of research close to reinforcement learning
- Objective is to learn **sequential decision processes**
- Online algorithms with **no training phase**
- Expert learning algorithms choose at each timestep which experts to follow
- Regret guarantees: finding the **best expert in sub-linear time**
- **Regret** $R_T = \sum_{t=1}^{T} f_t(a_t, y_t) - \inf_{e \in \mathrm{E}} \sum_{t=1}^{T} f_t\left(a_{e,t}, y_t\right)$

Expert interaction scheme

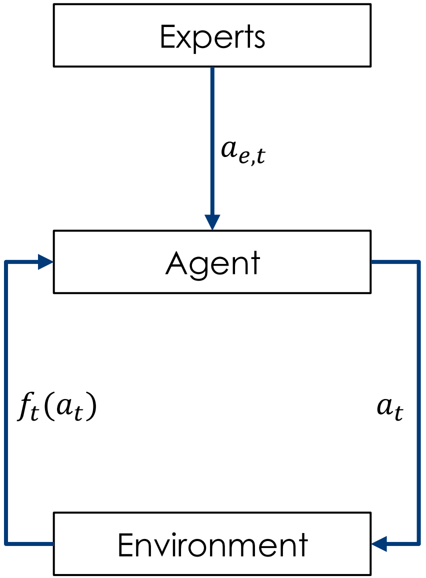

Expert Learning (2/3)

Combining strategies using Expert Learning

Cumulative % P&L on 2019

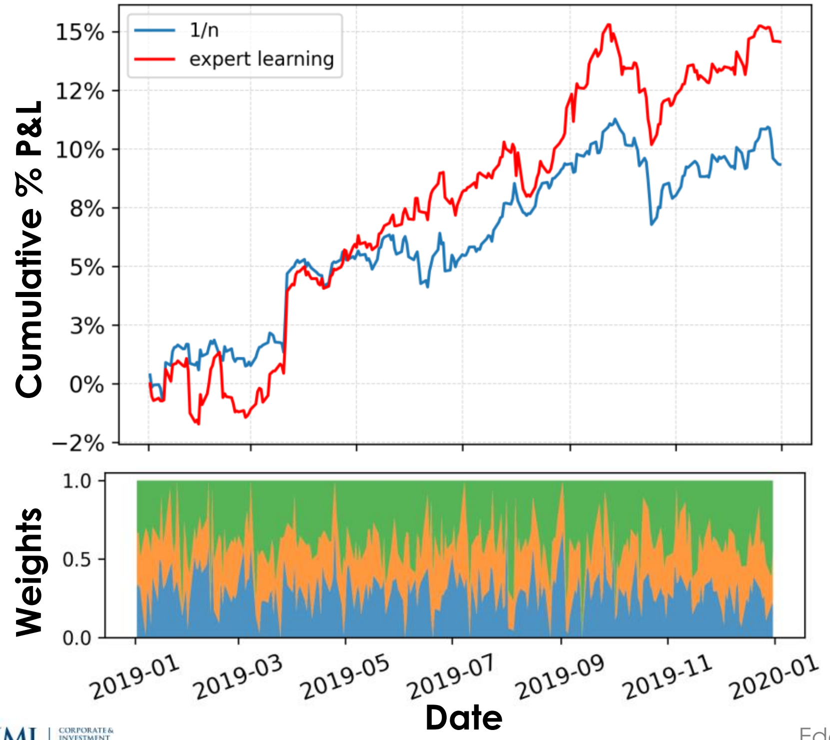

Performance metrics

	Sharpe	Vol
Expert learning	2.0	0.06
1/n	1.6	0.06

Comparing with the average helps to understand whether there is value in applying expert learning for dynamic rebalancing. Since volatility is below our target of 0.1, we can increase the size applied to the strategy

Expert Learning (3/3)

References on Expert Learning

- Wang, Yibo, et al. *Non-stationary projection-free online learning with dynamic and adaptive regret guarantees.* Proceedings of the AAAI Conference on Artificial Intelligence. 2024.
- Lu, Zhou, and Elad Hazan. "On the computational efficiency of adaptive and dynamic regret minimization." arXiv. 2022.
- Zhang, Lijun, Shiyin Lu, and Tianbao Yang. *Minimizing dynamic regret and adaptive regret simultaneously." International Conference on Artificial Intelligence and Statistic*. PMLR, 2020.
- Zhang, Lijun, Shiyin Lu, and Zhi-Hua Zhou. *Adaptive online learning in dynamic environments.* Advances in neural information processing systems. 2018.
- Luo, Haipeng, and Robert E. Schapire. *Achieving all with no parameters: Adanormalhedge*. Conference on Learning Theory. PMLR, 2015.
- Li, Bin, and Steven CH Hoi. *Online portfolio selection: A survey.* ACM Computing Surveys. 2014. Cesa-Bianchi, Nicolo, and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- Agarwal, A., Hazan, E., Kale, S., and Schapire, R. *Algorithms for portfolio management based on the newton method*. ICML, 2006.
- Zinkevich, Martin. *Online convex programming and generalized infinitesimal gradient ascent*. Proceedings of the 20th international conference on machine learning. 2003.
- Helmbold, David, et al. *Online portfolio selection using multiplicative updates*. Mathematical Finance, 1998.
- Cover, T. and Ordentlich, E. *Universal portfolios with side information*. IEEE Transactions on Information Theory, 1996.

Industry standards of quant trading strategies

Higher frequencies mean higher infrastructure costs

The focus of the tutorial is at this frequency

Frequency	Monthly	Daily	Intraday	Second	Nanosecond
Capacity	Billions	100s millions	10s millions	Few millions	Small size
Sharpe	~0.5	<=1	>=1	>=2	>8
Data availability	Company earnings	Price Sentiment Macro	Price	LOB TAQ	Tick by tick order flow

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

How can we improve?

Recap of current approach

Devise uncorrelated strategies

- Daily momentum strategy
- Intraday mean reverting
- Intraday seasonality

Normalize

Normalize the trading strategies

Combine

Combine using expert learning

How can we use machine learning to improve the underlying strategies?

- consider costs when generating the strategy
- move on from a strictly defined trading rule

Today's Agenda

Quantitative Trading

- Introduction
- Daily momentum
- Intraday mean reverting
- Intraday seasonality
- Strategy combination

RL for Quant Trading

- Reinforcement Learning for quantitative trading

Trading Futures with ML

- The Limit Order Book
- Feature engineering
- Strategy generation
- Optimizing execution

Reinforcement Learning for Quantitative 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:** price window, bid-ask spread, current portfolio, date/time
- **Action:** long, short, flat
- **Reward:** P&L – transaction costs

MDP Graphic

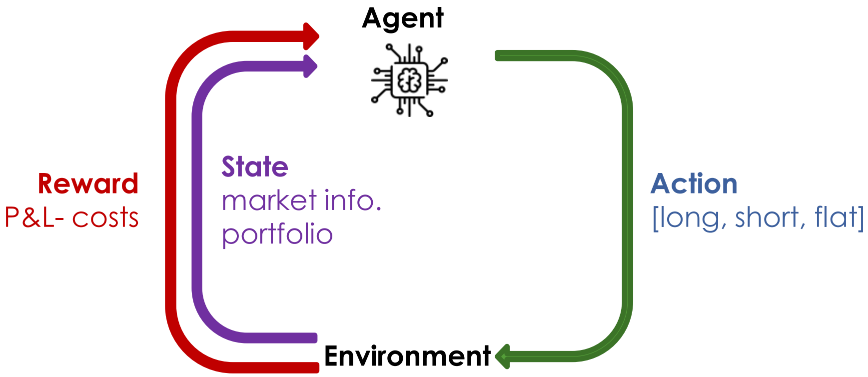

- The objective is finding the policy π which maximizes the discounted sum of the rewards
- $J_\pi = \mathbb{E}_\pi[\sum \gamma^t R_t]$

Reinforcement Learning for FX trading (1/3)

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

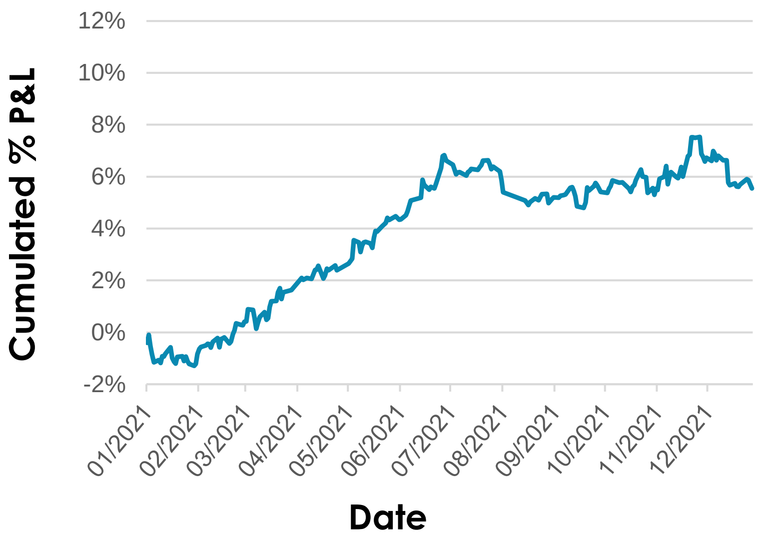

Performance Metrics

- 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.

Reinforcement Learning for FX trading (2/3)

Experimental results - policy

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

Actions chosen by the agent

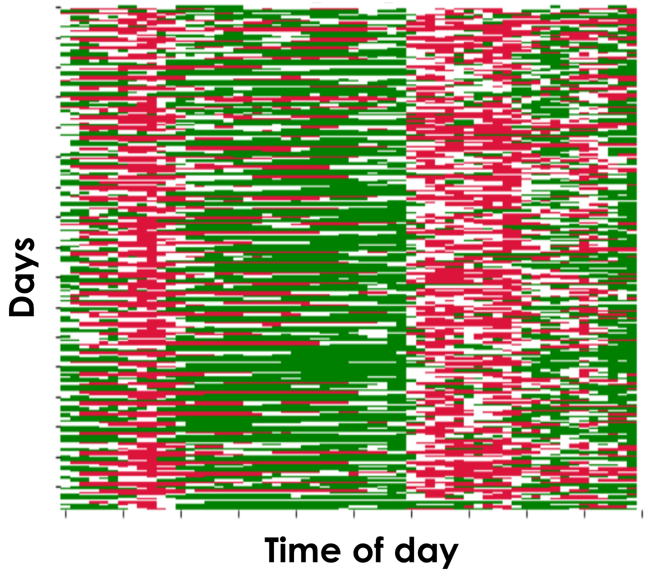

Riva, Antonio, et al. "Learning FX trading strategies with FQI and persistent actions." *ICAIF* 2021.

Reinforcement Learning for FX trading (3/3)

Experimental results - performance

P&L of backtest of RL strategies on 2021

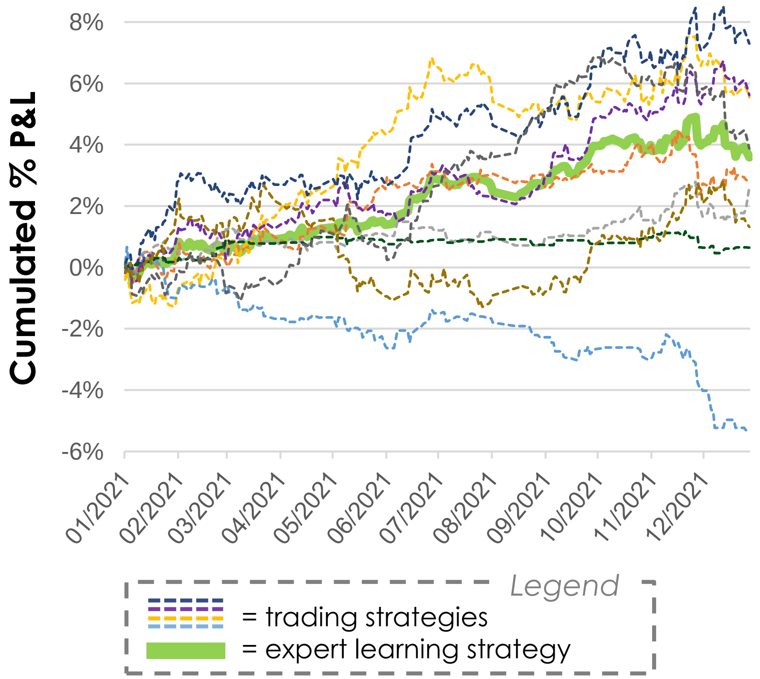

Performance metrics

Strategy	Return	Sharpe
Expert 1	6.8%	2.2
Expert 2	8.4%	2.1
Agent	4.3%	1.9
Expert 3	7.0%	1.8
Expert 4	3.3%	1.5
Expert 5	4.4%	1.3
Expert 6	3.6%	1.1
Expert 7	0.8%	0.7
Expert 8	1.7%	0.5
Expert 9	-6.4%	-2.5

Riva, Antonio, et al. "Addressing non-stationarity in FX trading with online model selection of offline RL experts." *ICAIF* 2022.

Today’s Agenda

Quantitative Trading

- Introduction
- Daily momentum
- Intraday mean reverting
- Intraday seasonality
- Strategy combination

RL for Quant Trading

- Reinforcement Learning for quantitative trading

Trading Futures with ML

- The Limit Order Book
- Feature engineering
- Strategy generation
- Optimizing execution

Intraday trading with Machine Learning

End-to-end workflow to build an intraday trading strategy

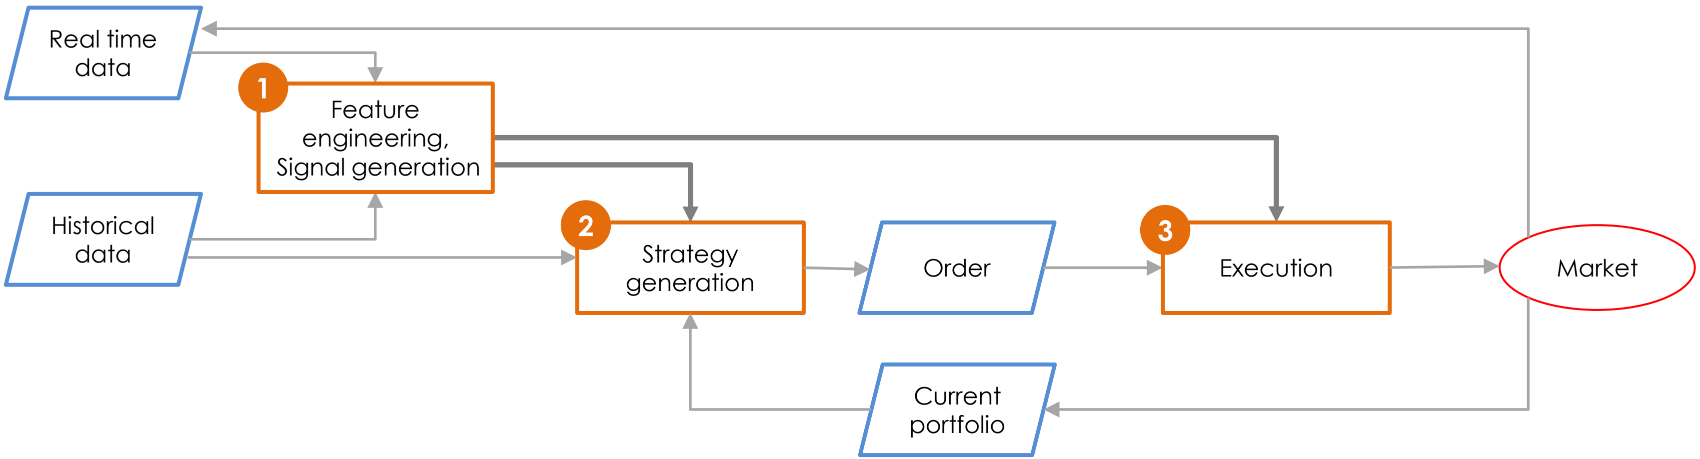

Automatic trading algorithms are comprised of 3 main components:

1. The **feature engineering** and **signal generation** component extracts **features** and **alpha signals** from the LOB
2. The **second component** takes the alpha signals and compressed LOB information to generate a **trading strategy**
3. The **execution algorithm** has the objective of minimizing transaction costs and controlling market impact

1

Limit Order Book (LOB) data

LOB data contains 10 price levels on the bid and on the ask

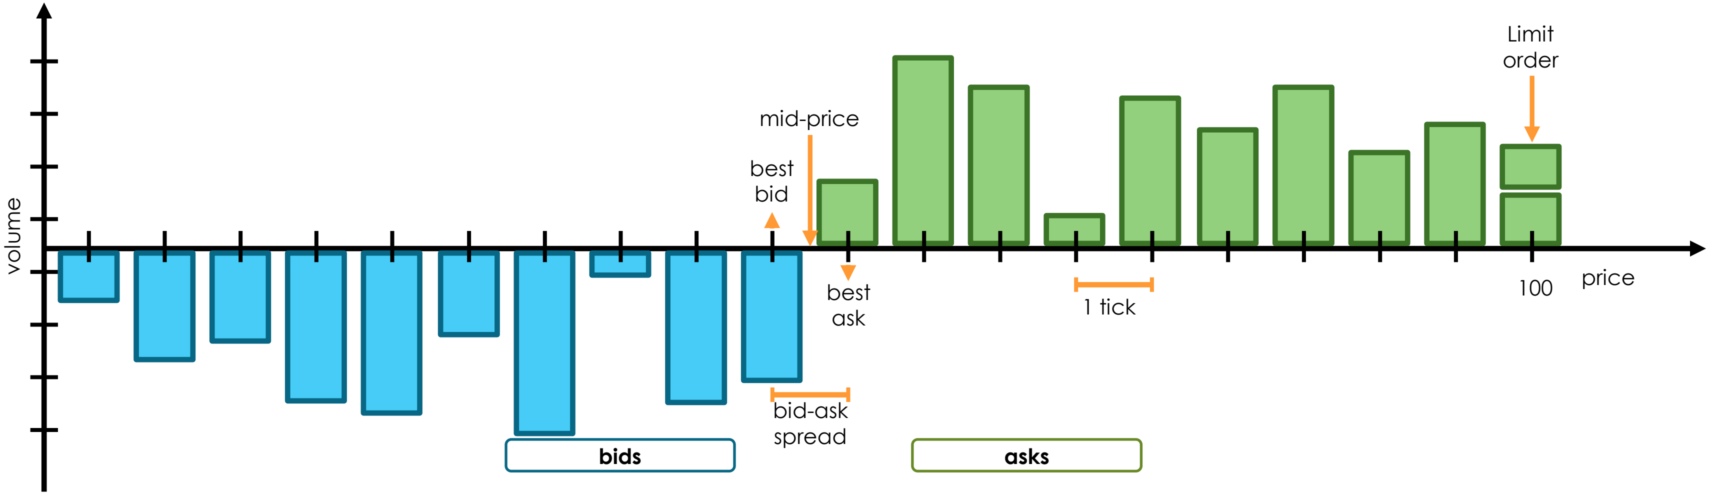

Trades and Quotes (TAQ) data

1

TAQ data contains only best bid, best ask and executed orders

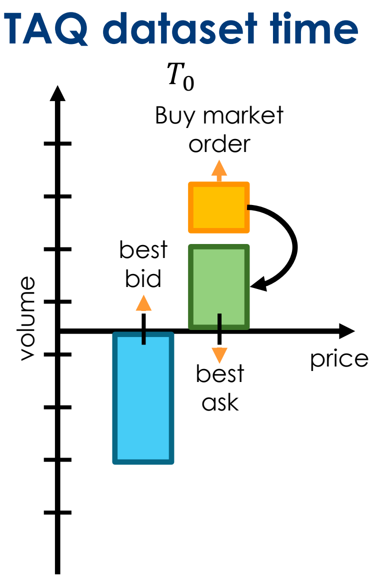

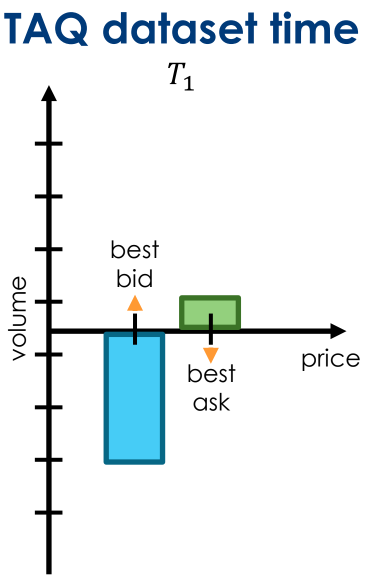

Processing the raw data

LOB and TAQ data have timestamps at nanosecond (10^{-9}) precision

Data Quality

- Handle nans/gaps (forward fill?)
- Merge the LOB and TAQ datasets
- Be careful with datetime features
- Handle roll dates

Working frequency

- At this step it is necessary to decide the frequency you want to work with:
  - Tick by tick
  - Downsample every x ticks
  - Downsample every x seconds

Extracting features from the LOB

We can use a classical approach, a machine learning approach or a combined one

Classical approach

Use hand crafted features such as:

- Autocorrelation of the price
- Order Flow Imbalance*

$$OF_{b,t}\begin{cases} v_{b,t} \text{ if } p_{b,t} > p_{b,t-1} \\ v_{b,t} - v_{b,t-1} \text{ if } p_{b,t} = p_{b,t-1} \\ -v_{b,t} \text{ if } p_{b,t} < p_{b,t-1} \end{cases} \qquad OF_{a,t}\begin{cases} -v_{a,t-1} \text{ if } p_{a,t} > p_{a,t-1} \\ v_{a,t} - v_{a,t-1} \text{ if } p_{a,t} = p_{a,t-1} \\ -v_{a,t} \text{ if } p_{a,t} < p_{a,t-1} \end{cases}$$

$$OFI_t = OF_{b,t} - OF_{a,t}$$

- Volume imbalance

$$\frac{v_b - v_a}{v_b + v_a}$$

- Trade imbalance

$$\sum_{N(t_{k-1})}^{N(t_k)} b_n - \sum_{N(t_{k-1})}^{N(t_k)} s_n$$

Machine Learning approach

- Use convolutional neural networks to extract features†
- Compress the information with autoencoders

Autoencoder example

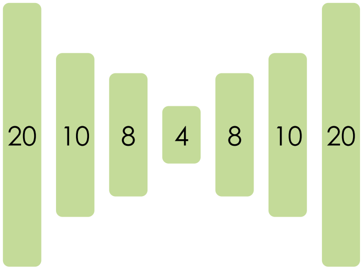

*Cont, Rama, Arseniy Kukanov, and Sasha Stoikov. "The price impact of order book events." *Journal of financial econometrics* 12.1 (2014)
†Zhang, Zihao, Stefan Zohren, and Stephen Roberts. "Deeplob: Deep convolutional neural networks for limit order books." (2018). *IEEE Transactions on Signal Processing* 67.11

1

Generating trading signals

The objective is to accurately predict the direction of the price movement

Defining the target

- target $\begin{cases} (\text{mid}_{t+xsec} - \text{mid}_t) < -\theta \rightarrow \text{Down} \\ -\theta < (\text{mid}_{t+xsec} - \text{mid}_t) < \theta \rightarrow \text{Stable} \\ (\text{mid}_{t+xsec} - \text{mid}_t) > \theta \rightarrow \text{Up} \end{cases}$
- $\text{mid}_{t+xticks} - \text{mid}_t$
- should I consider an average price to smooth out the noise?
- should I look at bid and ask prices instead of mid

Example of confusion matrix

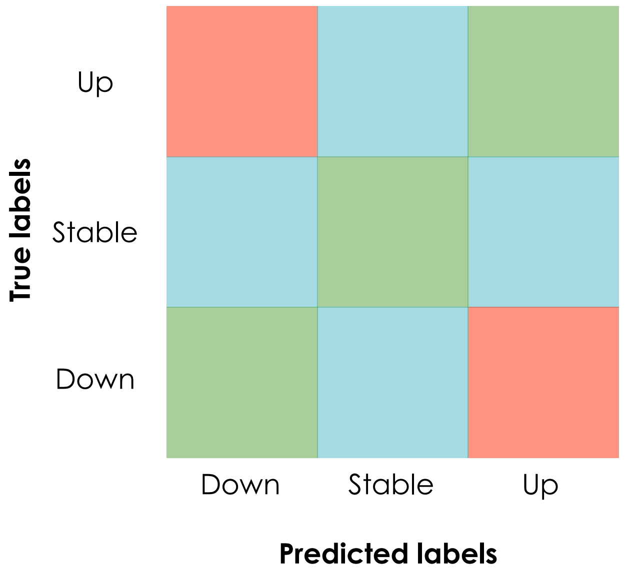

Classifier choice

- DeepLOB, LSTM, MLP
- Xgboost, Light GBM, Extra Trees
- Ensemble of predictions
- Hyperparameter tuning

References for LOB forecasting

1

- Briola, Antonio, Silvia Bartolucci, and Tomaso Aste. "*Deep Limit Order Book Forecasting*." *arXiv*. 2024
- Lucchese, Lorenzo, Mikko S. Pakkanen, and Almut ED Veraart. "*The short-term predictability of returns in order book markets: a deep learning perspective*." International Journal of Forecasting. 2024
- Kolm, Petter N., Jeremy Turiel, and Nicholas Westray. "*Deep order flow imbalance: Extracting alpha at multiple horizons from the limit order book*." Mathematical Finance. 2023
- Aït-Sahalia, Yacine, et al. *How and When are High-Frequency Stock Returns Predictable?*. National Bureau of Economic Research. 2022
- Briola, Antonio, et al. "*Deep reinforcement learning for active high frequency trading*." arXiv. *2021*
- Sirignano, Justin, and Rama Cont. "*Universal features of price formation in financial markets: perspectives from deep learning*." Machine Learning and AI in Finance 2021.
- Zhang, Zihao, Stefan Zohren, and Stephen Roberts. "*Deeplob: Deep convolutional neural networks for limit order books*." *IEEE Transactions on Signal Processing*. 2020
- Cont, Rama, Arseniy Kukanov, and Sasha Stoikov. "*The price impact of order book events*." *Journal of financial econometrics*. 2014

Defining the trading strategy using RL

RL has the task of optimizing the trading strategies taking into account transaction costs

State: trading signals, market information and current position
Action: can be the portfolio position or the trade
Reward: p&l – transaction costs

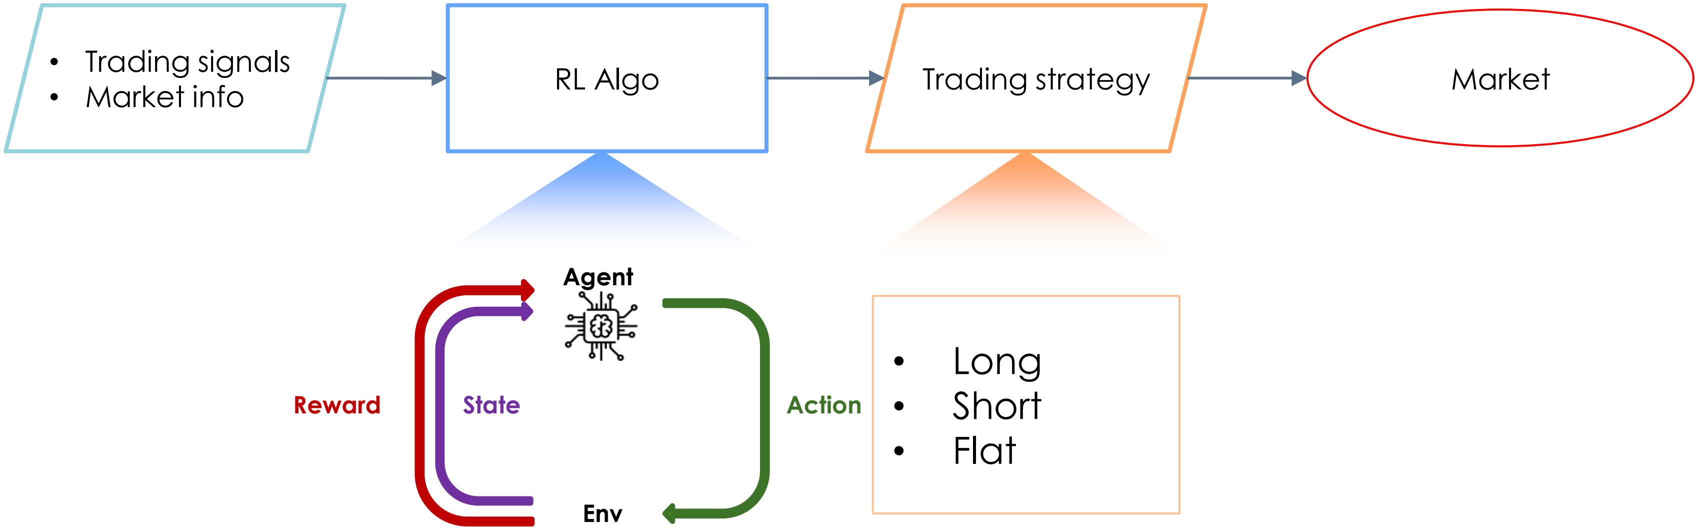

References on trading using RL

2

- Sun, Shuo, et al. "*TradeMaster: a holistic quantitative trading platform empowered by reinforcement learning*." Advances in Neural Information Processing Systems. 2024
- Sun, Shuo, Rundong Wang, and Bo An. "*Reinforcement learning for quantitative trading*." ACM Transactions on Intelligent Systems and Technology. 2023
- Brini, Alessio, and Daniele Tantari. "*Deep reinforcement trading with predictable returns*." Physica A: Statistical Mechanics and its Applications. 2023
- Riva, Antonio, et al. "*Addressing non-stationarity in FX trading with online model selection of offline rl experts*." ICAIF. 2022.
- Riva, Antonio, et al. "*Learning FX trading strategies with FQI and persistent actions.*" ICAIF 2021.
- Briola, Antonio, et al. "*Deep reinforcement learning for active high frequency trading*." arXiv. 2021
- Théate, Thibaut, and Damien Ernst. "*An application of deep reinforcement learning to algorithmic trading*." Expert Systems with Applications. 2021
- Carapuço, João, Rui Neves, and Nuno Horta. "*Reinforcement learning applied to Forex trading*." Applied Soft Computing. 2018
- Deng, Yue, et al. "*Deep direct reinforcement learning for financial signal representation and trading*." IEEE transactions on neural networks and learning systems. 2016
- Dempster, Michael AH, and Vasco Leemans. "*An automated FX trading system using adaptive reinforcement learning*." Expert systems with applications. 2006

Optimizing strategy execution

Optimizing execution can reduce transaction costs and market impact

Opportunity cost: execute immediately with high market impact or execute in time with the risk of a market movement?
Trading signal: shorter term compared to the one used in the trading strategy

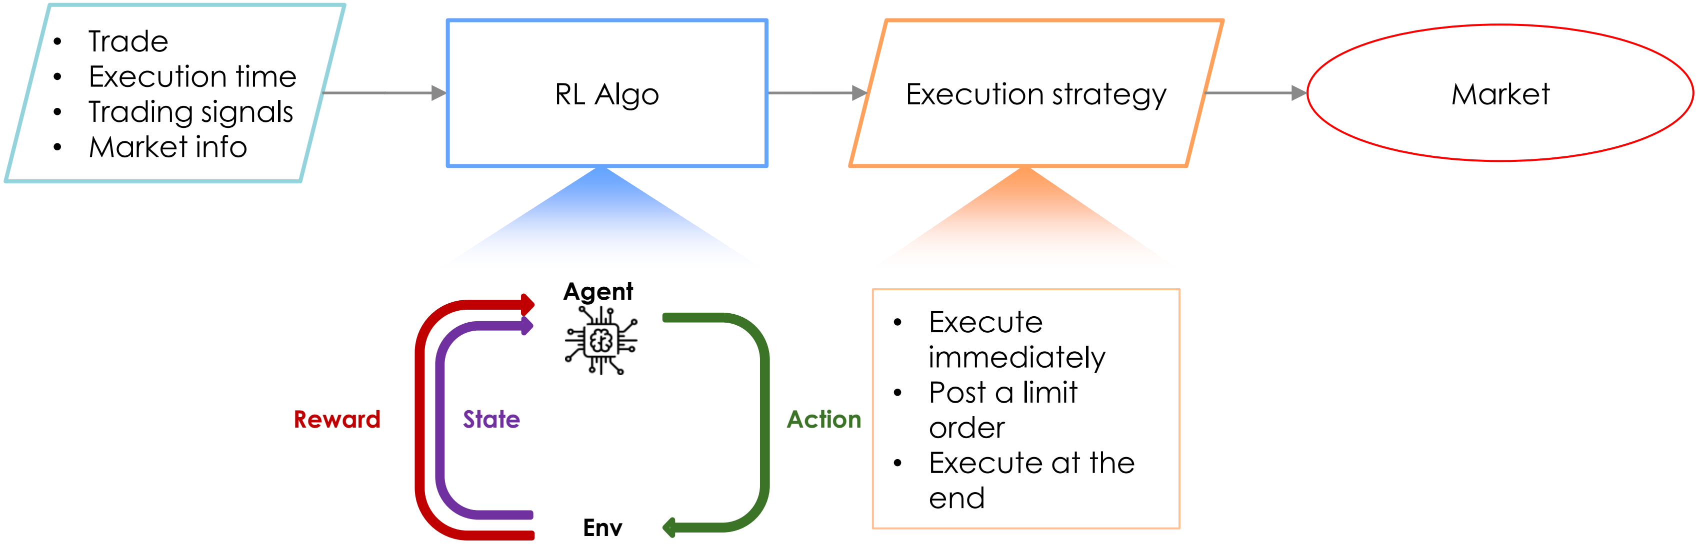

References on optimal execution

- Tommaso Mariotti, Fabrizio Lillo, and Giacomo Toscano. *From zero-intelligence to queue-reactive: limit-order-book modeling for high-frequency volatility estimation and optimal execution*. Quantitative Finance. 2023
- Marina Giacinto, Claudio Tebaldi, and Tai-Ho Wang. 2022. *Optimal order execution under price impact: a hybrid model. Annals of Operations Research*. 2022
- Siyu Lin and Peter A. Beling. *An End-to-End Optimal Trade Execution Framework based on Proximal Policy Optimization.* In IJCAI. 2020
- Michaël Karpe, Jin Fang, Zhongyao Ma, and Chen Wang. 2020. *Multi-agent reinforcement learning in a realistic limit order book market simulation*. ICAIF. 2020
- Brian Ning, Franco Ho Ting Lin, and Sebastian Jaimungal. 2018. *Double deep q-learning for optimal execution*. arXiv. 2018
- Álvaro Cartea, Sebastian Jaimungal, and Jason Ricci. 2018. *Algorithmic Trading, Stochastic Control, and Mutually Exciting Processes*. SIAM. 2018
- Dieter Hendricks and Diane Wilcox. 2014. *A reinforcement learning extension to the Almgren-Chriss framework for optimal trade execution*. IEEE Conference on Computational Intelligence for Financial Engineering & Economics. 2014
- Robert Almgren. *Optimal Trading with Stochastic Liquidity and Volatility*. SIAM Journal on Financial Mathematics 3. 2012

Intraday trading with Machine Learning

A generic workflow to generate a trading strategy with any listed asset

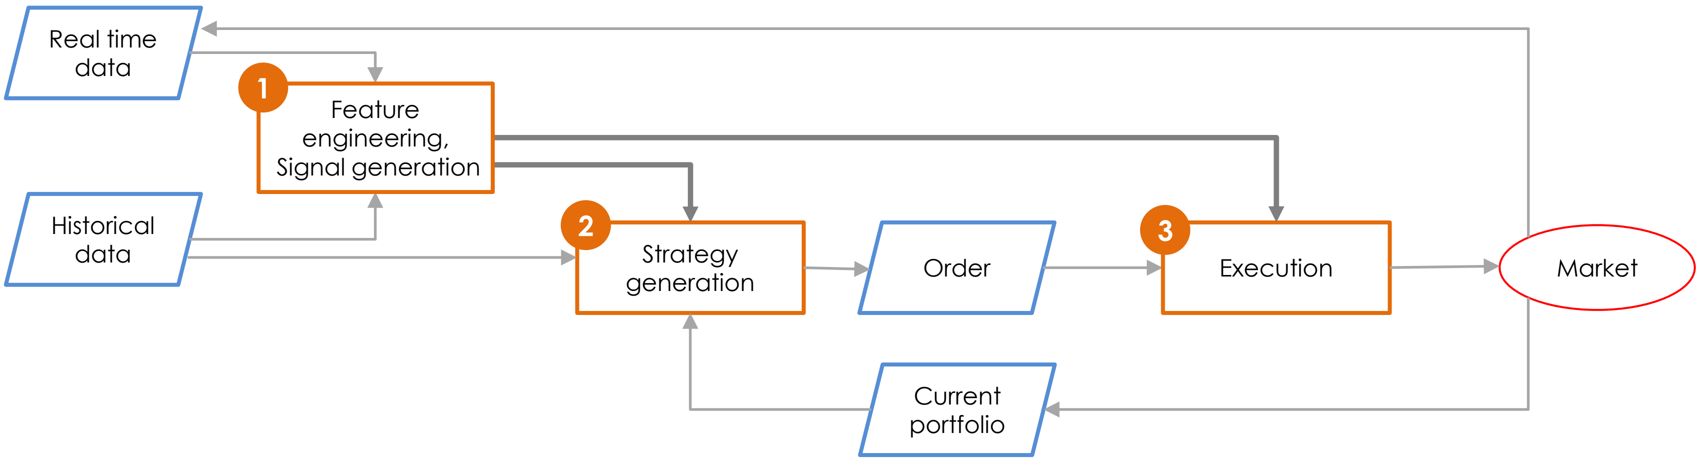

Automatic trading algorithms are comprised of 3 main components:

1. The **feature engineering** and **signal generation** component extracts **features** and **alpha signals** from the LOB
2. The **second component** takes the alpha signals and compressed LOB information to generate a **trading strategy**
3. The **execution algorithm** has the objective of minimizing transaction costs and controlling market impact

Conclusions

Quantitative Trading

- Daily momentum
- Intraday mean reverting
- Intraday seasonality
- Strategy combination

RL for Quant Trading

- Reinforcement Learning for quantitative trading

Trading Futures with ML

- The Limit Order Book
- Feature engineering
- Strategy generation
- Optimizing execution

Contact info

Edoardo Vittori
edoardo.vittori@intesasanpaolo.com