Mean Reversion and Ornstein-Uhlembeck (OU) processes

pricing | risk mgmt | exec. | automation

Reference:
Review of Statistical Arbitrage, Cointegration and Multivariate Ornstein-Uhlenbeck, Meucci (2009)
In the one-dimensional case

$$dX_t = -\theta(X_t - \mu)dt + \sigma dB_t \tag{15}$$

which admits the analytical solution

AR(1)

$$X_{t+\tau} = (1 - e^{-\theta\tau})\mu + e^{-\theta\tau}X_t + \epsilon_{t,\tau} \tag{16}$$

where
$\epsilon_{t,\tau} \equiv \int_t^{t+\tau} e^{\theta(u-\tau-t)}\sigma dB_u \sim N(0, \sigma_\tau^2)$

$\sigma_\tau^2 = \int_t^{t+\tau} \left(e^{\theta(u-\tau-t)}\right)^2 \sigma^2 du$

$-\vartheta\tau = \ln\frac{1}{2}$ $\quad \vartheta\tau = \ln 2$ $\quad \tau = \frac{\ln 2}{\vartheta}$

Let's define $e^{-\vartheta\tau} = \frac{1}{2}$ $\quad = \mu + e^{-\vartheta\tau}(X_t - \mu)$

$$x_{t+\tau} \equiv E(X_{t+\tau}|X_t) = (1 - e^{-\theta\tau})\mu + e^{-\theta\tau}X_t \tag{17}$$

$$\sigma_\tau^2 \equiv \mathrm{Var}(X_{t+\tau}|X_t) = \frac{\sigma^2}{2\theta}\left(1 - e^{-2\theta\tau}\right) \tag{18}$$

then the conditional probability distribution is given by

$$\mathrm{pdf}(X_{t+\tau}|X_t) \sim N(x_{t+\tau}, \sigma_\tau^2) \tag{19}$$

as τ grows to infinity

$$\boxed{\mathrm{pdf}(X_\infty) \sim N(\mu, \sigma^2/2\theta)} \tag{20}$$

StatArb and Cointegration

A multivariate process is cointegrated if there exists a linear combination of its entries which is stationary.
Practitioners usually look for a weaker condition, i.e. a linear combination which, based on historical data, behaves as a mean reverting OU-process with $\theta^{-1} \ll T$, where T is their typical time horizon for closing a trade.

Search for cointegration: PCA

$X \begin{bmatrix} P_1 & \cdots \\ P_2 & \\ P_3 & \\ \vdots & \end{bmatrix}^T \quad X^T = \begin{bmatrix} \quad \end{bmatrix}$

Principal Component Analysis (PCA) is often used as a data driven approach. Given a multivariate time series X^T where each column of X represents a single realization of N financial securities (i.e. $X_{i,t}$ is the price of security i at time index t), usually but not necessarily selected based on macroeconomic criteria, the following computational steps are applied:

- Estimate the covariance matrix of the time series $\Sigma = \mathbb{E}(XX^T) - \mathbb{E}(X)\mathbb{E}(X^T)$
- Compute its eigenvalues and eigenvectors, such that $\Sigma = \mathrm{A}\Lambda\mathrm{A}^T$, $\mathrm{A}\mathrm{A}^T = I$

where A is the matrix whose columns are the eigenvectors $\{e_i\}_1^N$ and Λ is a diagonal matrix whose elements are the positive eigenvalues ordered such that $\lambda_1 \geq \lambda_2 ... \geq \lambda_N$ and $\Lambda_{i,j} = \lambda_i \delta_{i,j}$.

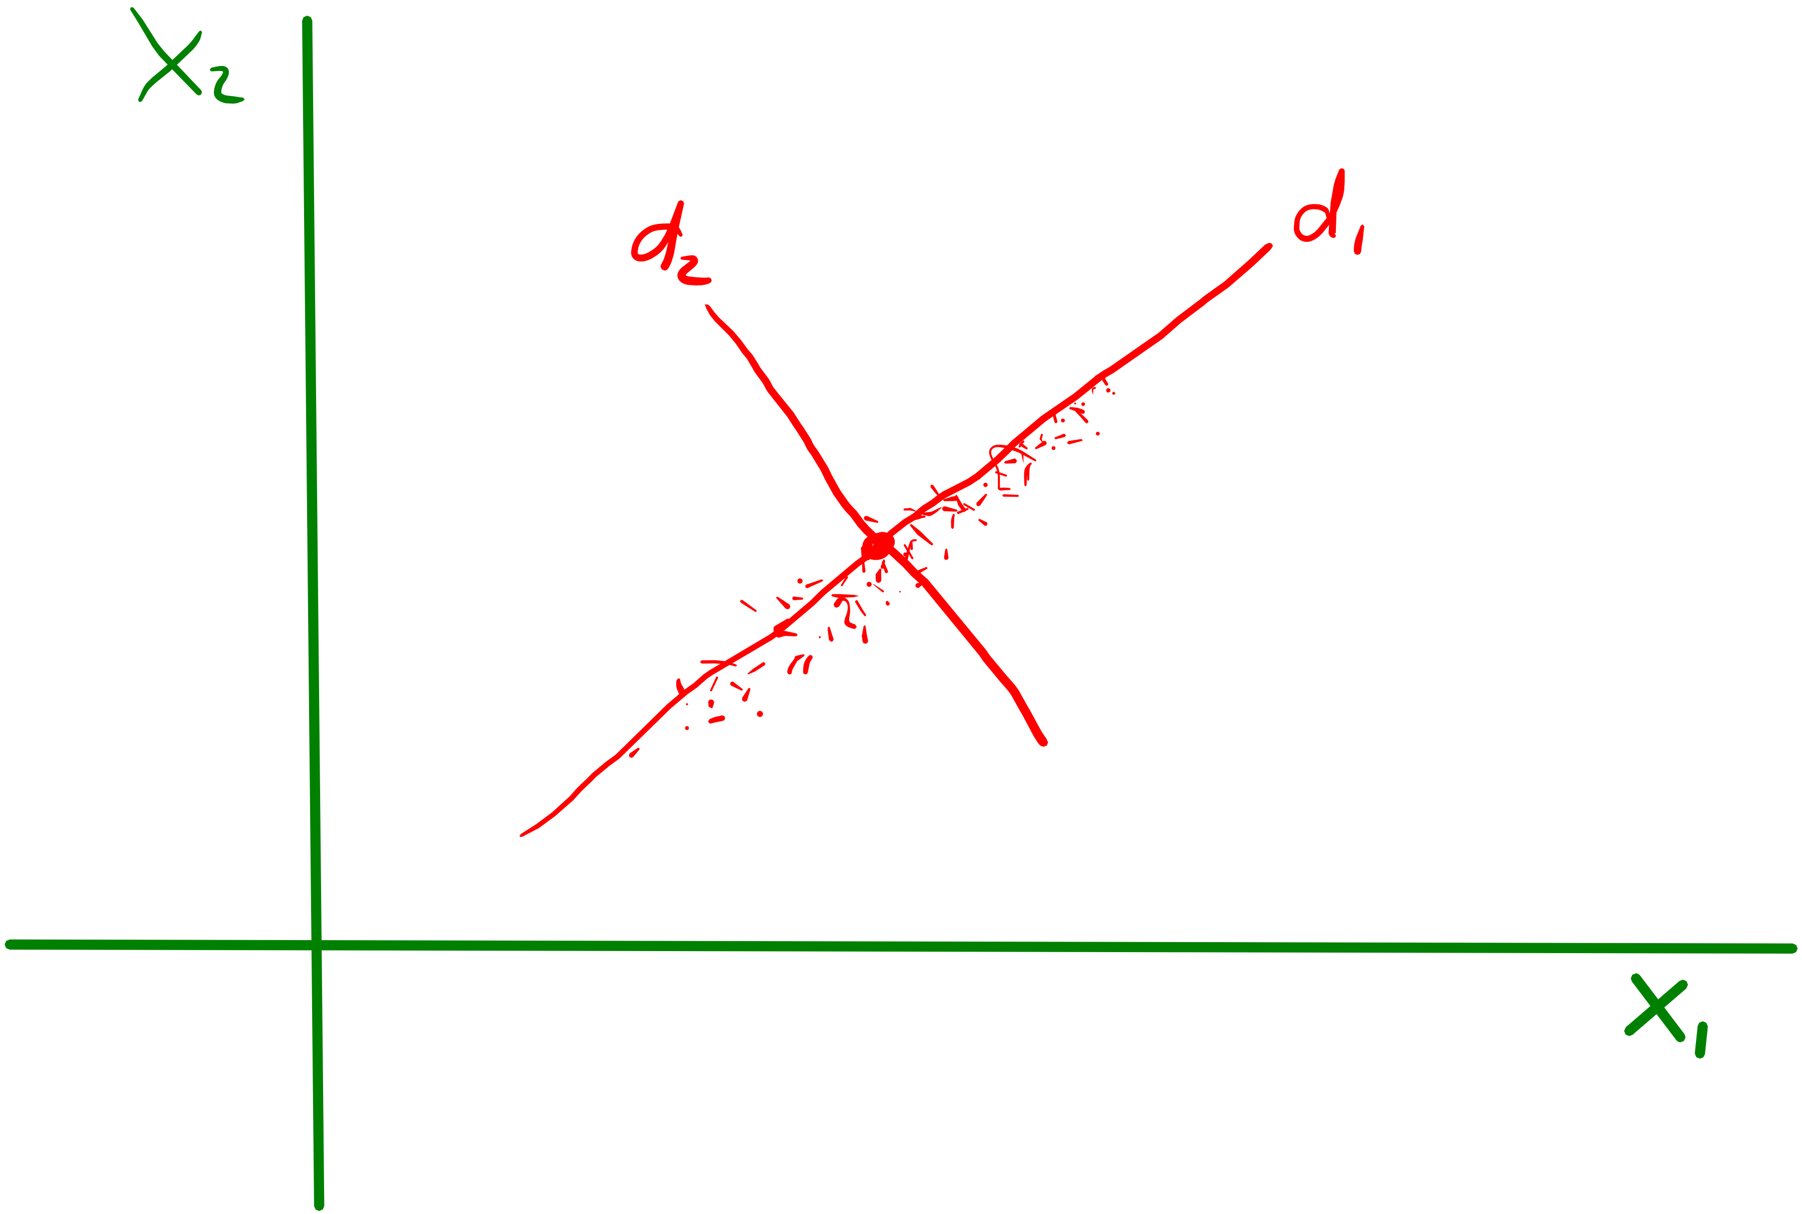

X_2
d_2
d_1
X_1

PCA

$X^T \qquad w \qquad X^T w$

$(X^T w)^T X^T w = w^T X X^T w$

$t_1, t_2, \ldots, T \qquad N > 1$

$w^T w = 1$

The eigenvector e_N relative to the smallest eigenvalue λ_N satisfies the following condition

$$e_N = \underset{e_i}{\operatorname{argmin}} \left(\operatorname{Var}\left(X^T e_i\right)\right) = \underset{\|w\|=1}{\operatorname{argmin}} \left(\operatorname{Var}\left(X^T w\right)\right) \tag{21}$$

As can be shown by noting that $\operatorname{Var}\left(X^T w\right) = w^T \Sigma w$ and taking into account that condition $\|w\| = 1$ implies that a stationary point for $w^T \Sigma w$ can be obtained only when w is an eigenvector. Finally $e_i^T \Sigma e_i = \lambda_i$ which is minimum for $i = N$.

$w^T \Sigma w - \lambda \cdot w^T w = \text{minimum} \qquad \Sigma e_i = \lambda_i e_i$

PCA - Considerations

- Unit length $\|w\|$ vs. usual optimal portfolio condition $\sum w_i = 1$
- Data are usually normalized after subtracting the mean (tricky)
- Financial markets are far from being stationary. Sensitivity on future covariance estimation error is significant.

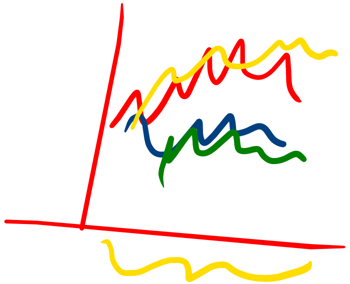

$$-X^T$$

$$\Sigma = Var(X^T)$$

$$\begin{bmatrix} | & & | \\ e_1 & \cdots & e_N \\ | & & | \end{bmatrix} \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_N \end{bmatrix}$$

$$X^T e_1 \xrightarrow{var} \lambda_1$$

$$\vdots \qquad \vdots$$

$$X^T e_N \xrightarrow{var} \lambda_N$$

$$\xrightarrow[AR(1)]{} \underline{\theta, \sigma^2, \mu} \rightarrow \frac{\sigma^2}{2\theta} \qquad \underline{\theta^{-1} \ll T}$$

$$\mu$$

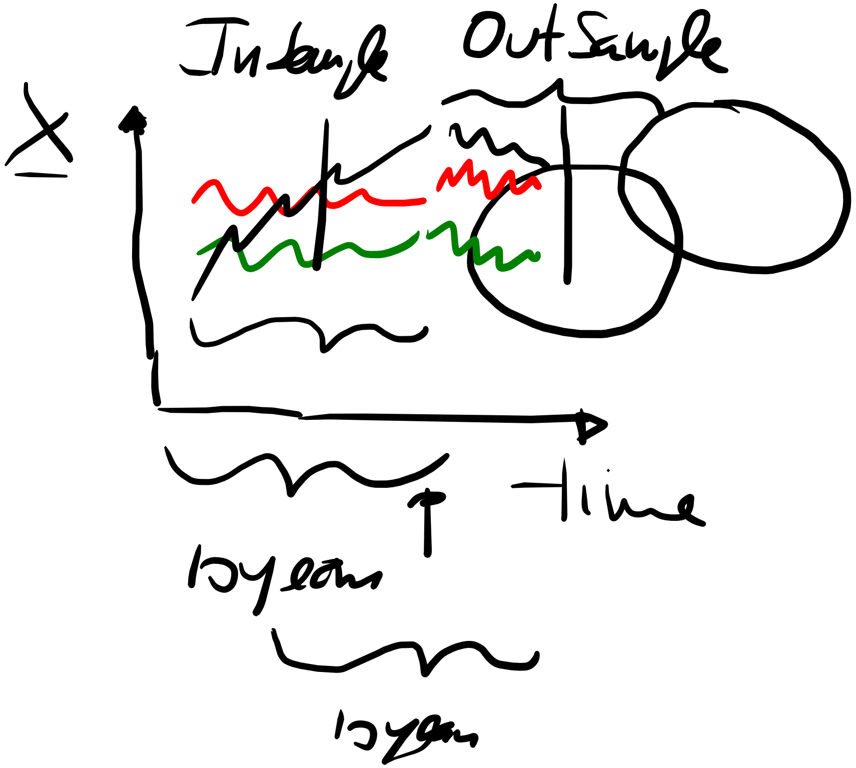

WARNING

DON'T USE
FUTURE INFORMA.
TO ESTIMATE PARAMETERS

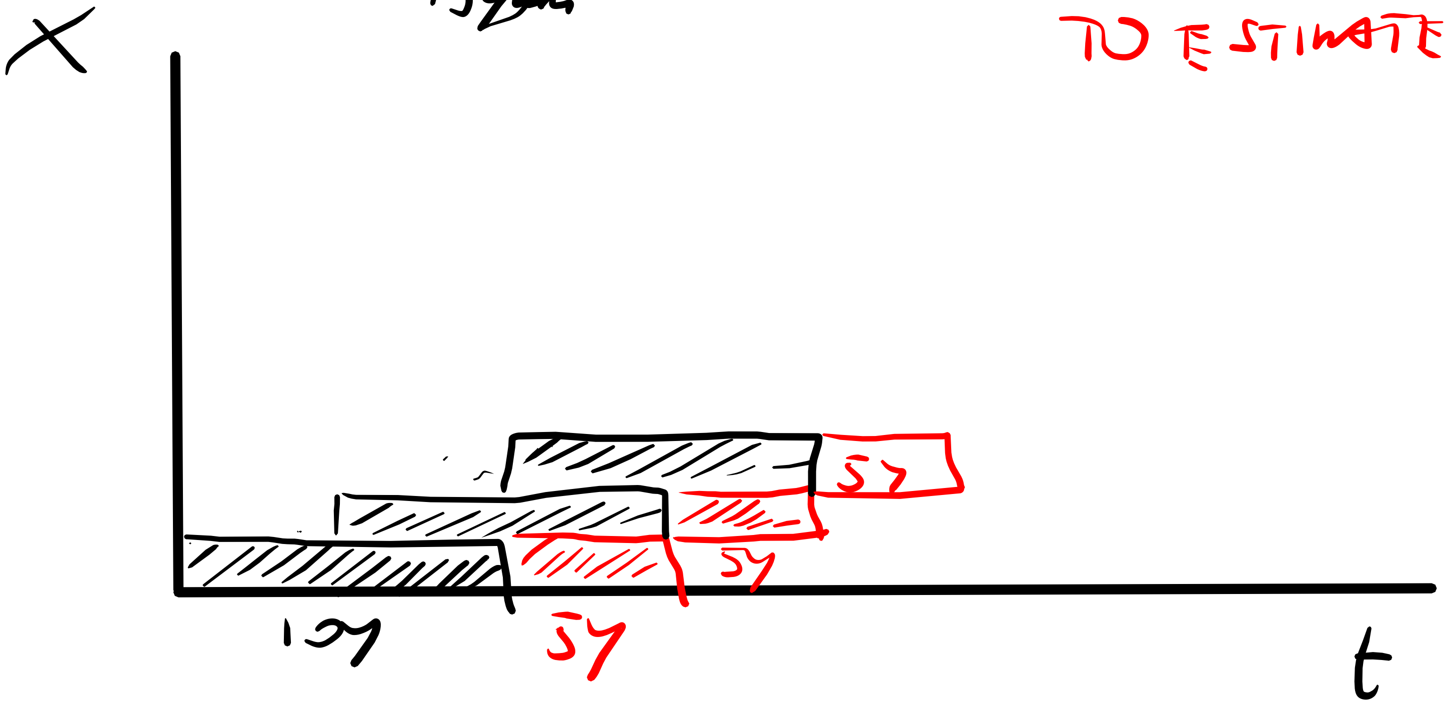

Fitting OU process to principal components

Ses Matlab scripts.
FitOU.m (multidimensional case, linear regression VAR(1))
FitOUAR1.m (one-dimensional, linear regression AR(1))
FitOUCheck.m (one-dimensional, Unconditional Moments Estimation Approach).

Caveats

- The smallest eigenvalues and eigenvectors are subject to high estimation error
- Execution costs not negligible compared to signal
- Thetas optimistic bias

See also:
Estimation of Ornstein-Uhlenbeck Process Using Ultra-High-Frequency data with Application to Intraday Pairs Trading Strategy, Holy and Tomanova, (2018), Appendix B, for an introduction and an extension to take into account microstructure effects modeled using additive Gaussian noise.
Mean-Reverting Spread Modeling: Caveats in Calibrating the OU Process, Hansen Pei, 2021.

Statistical Arbitrage: z-score

A common approach is to scale a trader position according to the so called z-score z_t defined as follows:

$$z_t = \frac{(X_t - \mu)}{\sqrt{\sigma^2/2\theta}} \tag{22}$$

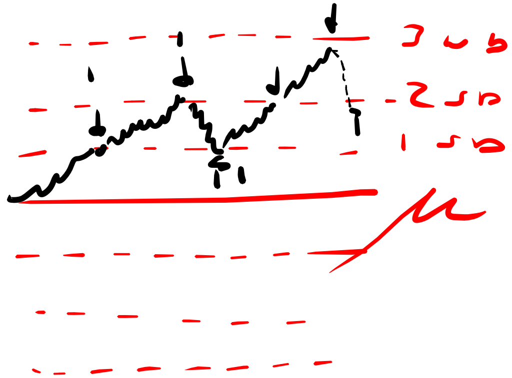

and open a position α_t according to

$$\alpha_t \propto -fix(z_t) \tag{23}$$

where *fix()* represents rounding towards zero (MATLAB syntax). Positions are typically closed or reduced after a significant mean reversion towards the mean occurs, for example after z_t reverts towards zero by 0.5 with respect to its value at the opening of a new position. Alternatively a position is cut down to zero when losses approach a given threshold (stop loss).

Implementing a strategy on a OU processes: theory

It can be shown (Boguslavsky and Boguslavskaya) that the optimal strategy for an agent under a log utility $U(W) = \log(W)$ and for a security whose price follows an OU-process

$$\alpha_t = -W_t(X_t - \mu)\frac{\theta}{\sigma^2} \tag{24}$$

where W_t is the agent's wealth and the other variables are defined as in eq. 15, and

$$dW_t = \alpha_t dX_t \tag{25}$$

Optimal execution: outline

L (BUY AMOUNT)

- Performance metrics (benchmarks)
- Execution Cost and Implementation Shortfall
- Limit vs. Market Orders
- Limit orders and Markouts
- Latency

Optimal Execution

Problem: optimize the weighted execution price given the quantity according to some constrains, typically the time required to complete a trade or the participation rate.
Example: buy 100 shares of Netflix Inc. at the lowest average price within one minute.
Benchmarks: arrival price (e.g. current market weighted ask price when the decision is made), VWAP (Volume Weighted Average Price), TWAP (Time Weighted Average Price), etc.

BUY
100 shares of Netflix

t	Q_{bid}	P_{bid}	P_{ask}	Q_{ask}
	50	256.7	$\boxed{258.9}$	50
			$\boxed{260.1}$	50

ARRIVAL PRICE

1. $0.5(256.7 + 258.9)$
2. $$\frac{50 \cdot 258.9 + 50 \cdot 260.1}{50 + 50}$$

$\boxed{VWAP}$ (DAILY) $\quad \{p_t\} \to \{q_t\}$

$$L_s = \frac{\sum_t q_t \cdot p_t}{\sum_t q_t}$$

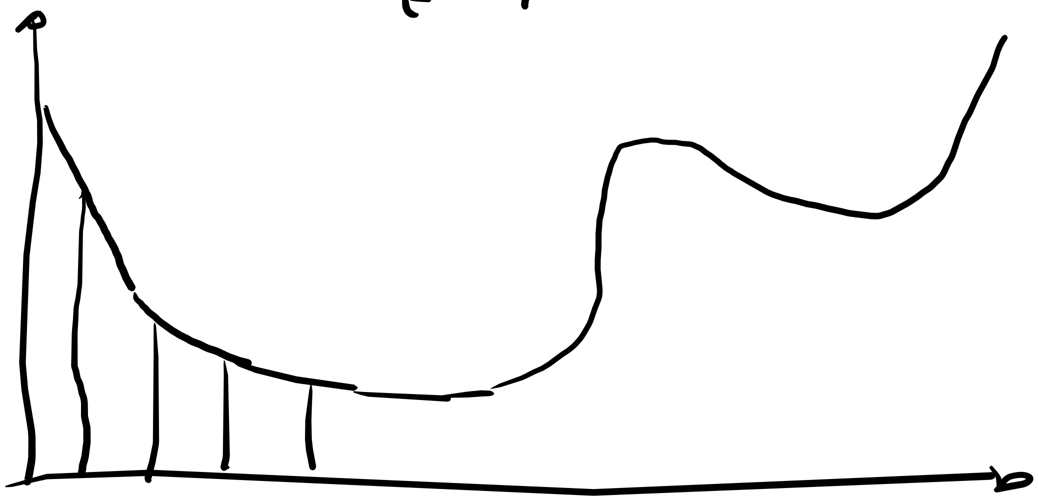

Optimal execution: execution cost

BUY   Ex cost = Pt − Parrival

weighted average price   VWAP TWAP

Also known as the 'slippage', execution cost refers to the difference between the benchmark price and the actual execution price. It may have a substantial negative impact on the profitability of a strategy.

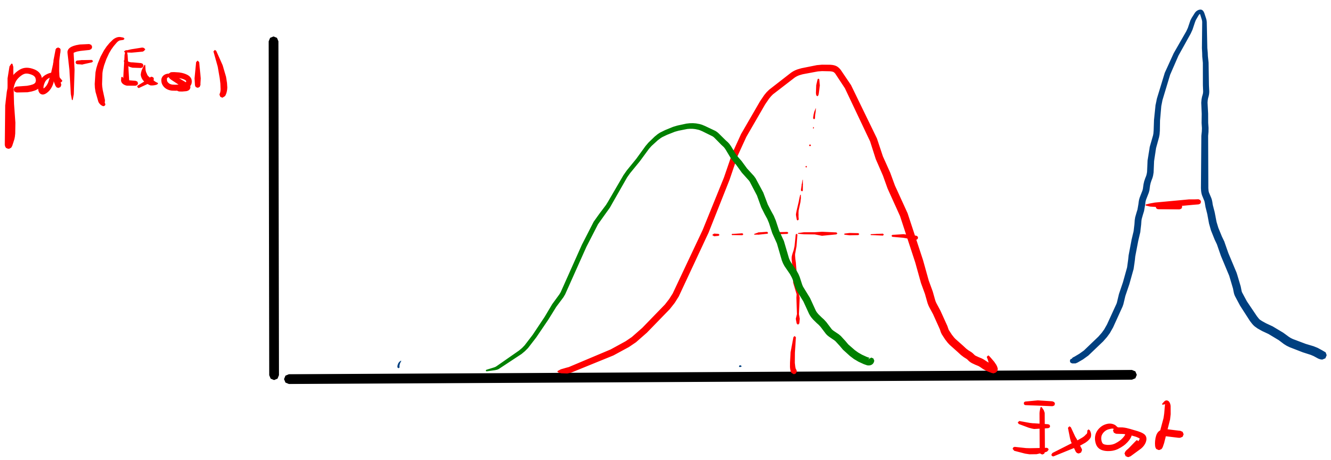

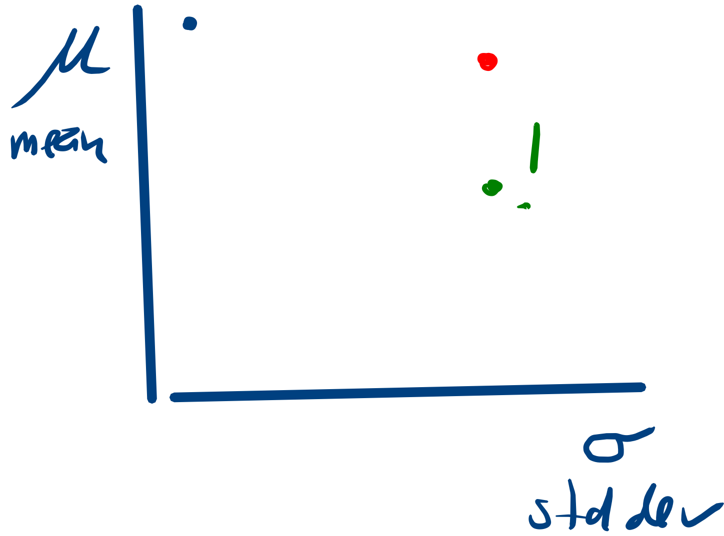

Optimal execution: Implementation Shortfall

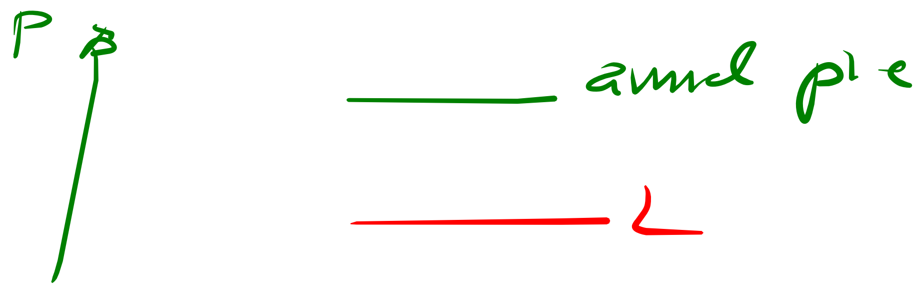

Implementation Shortfall = Execution Cost + Opportunity Cost
Execution Cost : actual traded prices vs. benchmark
Opportunity Cost : final portfolio different from desired portfolio
To be understood in a probabilistic setting (expected cost vs risk).

Implementation Shortfall

Given the initial portfolio quantities n_0 and prices π_0, as well as the desired portfolio quantities v and terminal prices π_1 compute the Implementation Shortfall (IS), taking into account execution prices p and actual final portfolio quantities n_1.

$$IS = (v - n_1)\pi_1 \tag{26}$$

It can be shown that

$$IS = \underbrace{(n_1 - n_0)(p - \pi_0)}_{\text{execution cost}} + \underbrace{(v - n_1)(\pi_1 - \pi_0)}_{\text{opportunity cost}} \tag{27}$$

using self financing conditions

$$n_0\pi_0 = v\pi_0 \tag{28}$$

$$n_0 p = n_1 p \tag{29}$$

A Reinforcement Learning Approach to Optimal Execution - Ref. Banca Sella Holding internal report

- Cost

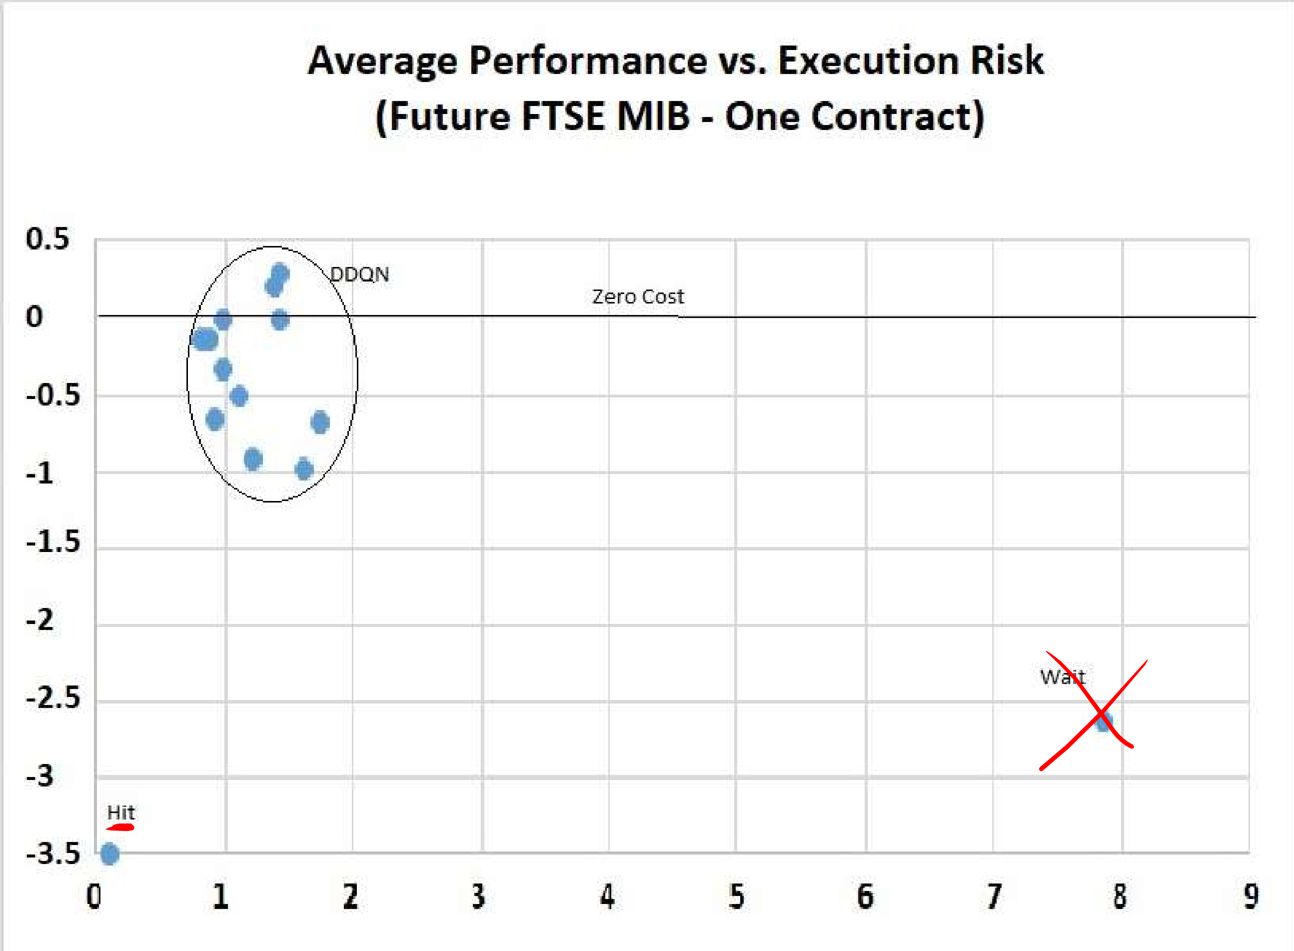

Reinf. learn.

lift immed (market order)

wait

$\boxed{\text{DDQN}}$

$Q(s, a)$

Limit vs market orders

$U(W) = -e^{-\gamma W}$

$\underline{n_s}$ — $\underline{n_s}$ 1° CASE

— $\underline{n_s + 1}$ 2° CASE $\underline{\text{BUY at } L}$

The task: choose between placing a limit order, a market order or do nothing when buying (selling) one share of a stock given the ask (bid) price and a model for the probability of being hit (lifted) by other market participants.
Let us assume a terminal wealth distribution $W \sim N(\mu_W, \sigma_W^2)$ and a utility framework (Constant Absolute Risk Aversion) $U(W) = -e^{-\gamma W}$.
W is computed as $\underline{n_s} * \underline{S}$ where n_s is the number of shares the trader holds after execution, and $\underline{S \sim N(\mu_{\cancel{X}}, \sigma_{\cancel{X}}^2)}$ is the shares terminal price distribution. If n is the initial number of shares, $n_s = n$, when no order is executed, while $n_s = n + 1$ for a completed buy order of one share, ($n_s = n - 1$ for a sell).

$S \sim N(\mu_x, \sigma_x^2) \rightarrow \underline{W \sim N\left(n_s \cdot \mu_s, n_s^2 \sigma_s^2\right)}$ 1° CASE

Execution probability

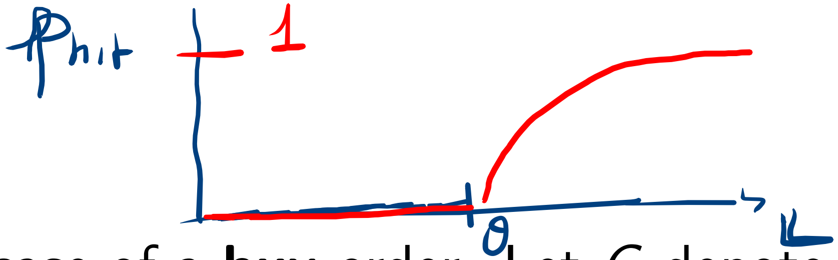

Consider the case of a **buy** order. Let C denote the minimum unexpressed reservation price across all sellers in the market willing to sell one share. A buy order at price L for one share will be executed only if $L \geq C$. Let's assume the following probability distribution for C for $C \geq \theta$ (zero otherwise)

$$\boxed{f(C) = \lambda e^{-\lambda(C-\theta)}} \tag{30}$$

The execution probability Pr_{Hit} of a buy order with limit price L is then given by

$$\boxed{Pr_{Hit}(L) = 1 - e^{-\lambda(L-\theta)}} \tag{31}$$

for $L \geq \theta$, and zero otherwise.

1° CASE

$n_s \to W \sim N(\mu_s \cdot n_s,\ n_s^2 \cdot \sigma_s^2)$

$\to EU_{safe}(W) = -e^{-\gamma\left(\mu_s n_s - \frac{1}{2}\gamma n_s^2 \sigma_s^2\right)}$

2° CASE

$n_s \to n_s + 1 \quad W \sim N\left(\mu_s \cdot (n_s+1) - L,\ (n_s+1)^2 \sigma_s^2\right)$

$\to EU_{hit}(L) = -e^{-\gamma\left(\mu_s (n_s+1) - L - \frac{1}{2}\gamma (n_s+1)^2 \sigma_s^2\right)}$

$EU(L) = p_{hit}(L) \cdot EU_{hit}(L) + (1 - p_{hit}(L)) \cdot EU_{safe}$

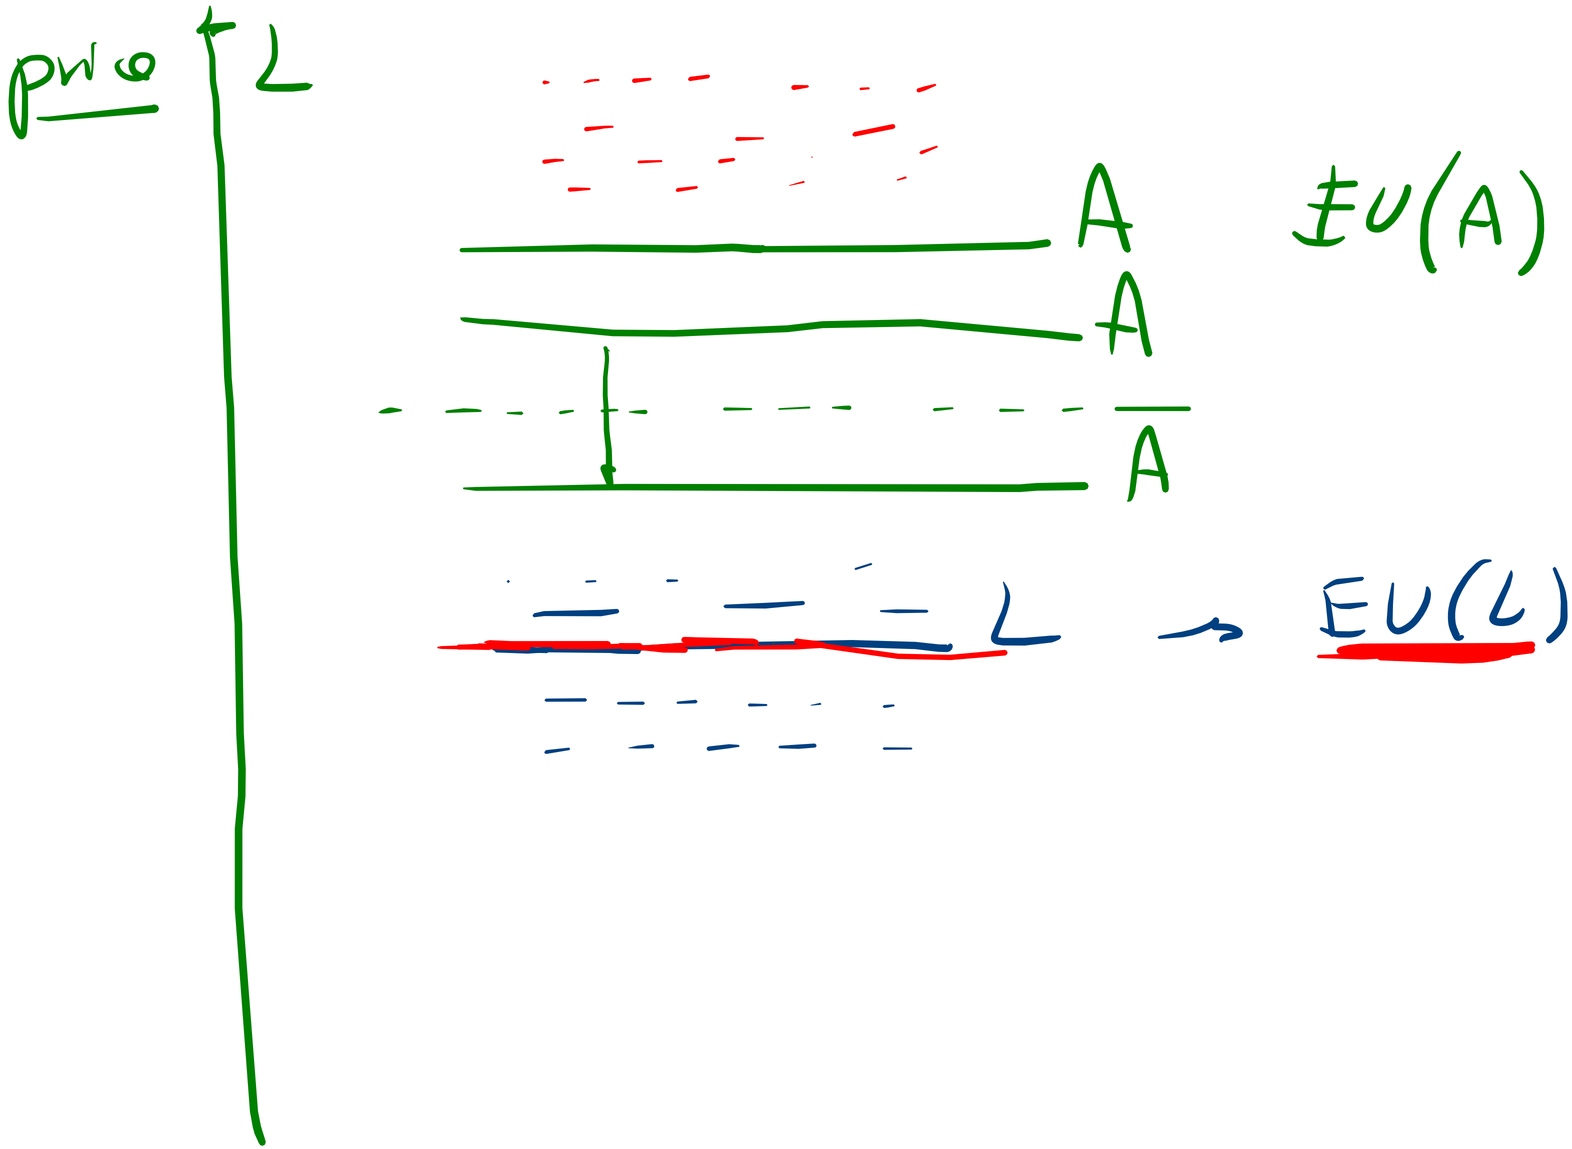
pro
L
A
A
A
L
EU(A)
EU(L)

Case: $\gamma = 1, \mu_X = 1, \sigma_X^2 = 1, \lambda = 1, \theta = 0$

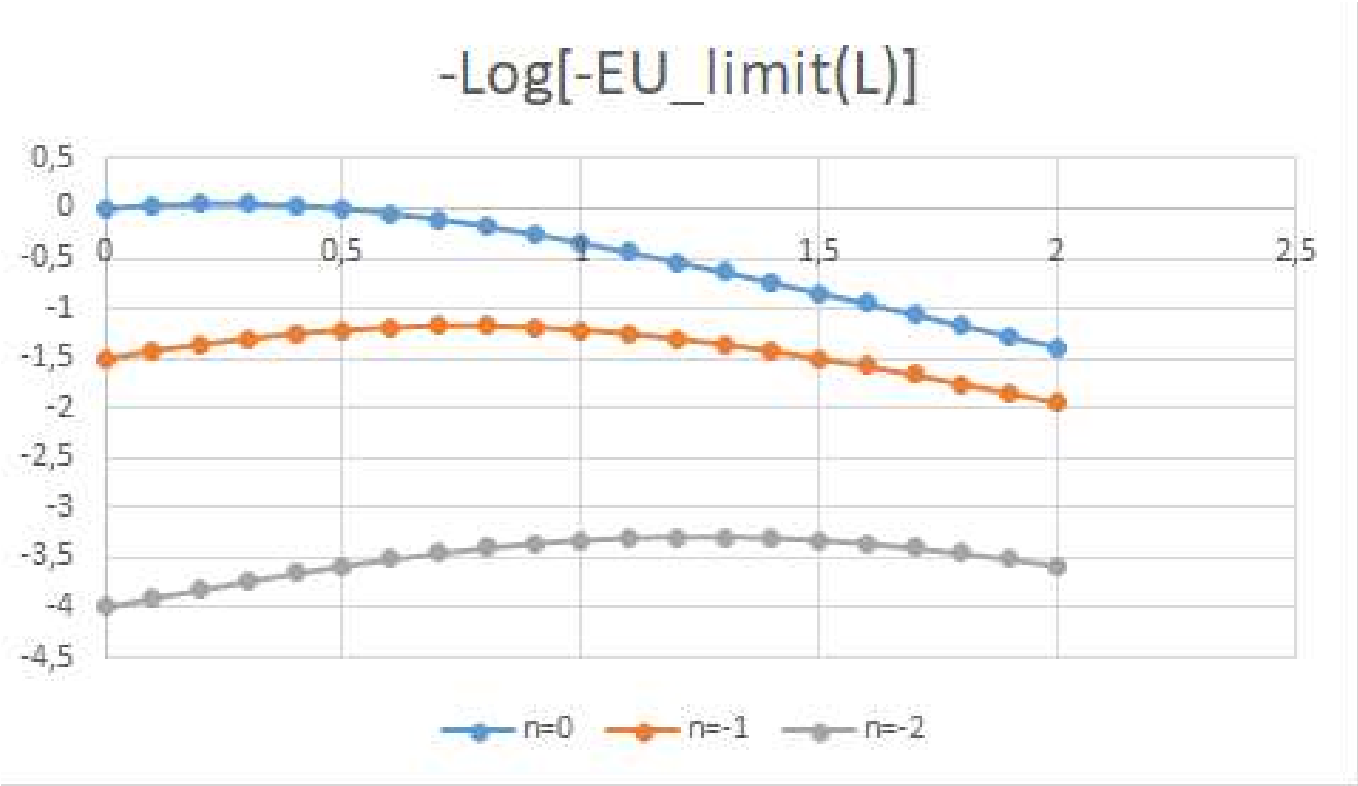

For $n = -1$, trader short one unit, optimal limit price is 0.75 with a 53% probability of being hit.

Immediate execution: gravitational pull

A market order to buy is preferable to a limit order when there is a favorable ask at price A such that

$$E(U_{MarketOrder,A}) > E(U_{LimitOrder,L}) \quad \forall L < A \tag{32}$$

Warning: our main assumption about the share's final price distribution is that $S \sim N(\mu_X, \sigma_X^2)$ does not depend on whether we do nothing or we place a limit order and it is executed or not. Moreover, how other market participants will react is not considered here.

Markouts

In order to check if our previous assumption holds, we can check empirically if our actions, i.e. placing a limit order and/or being executed, change our security terminal price.
In particular, $E(S_{noexe} - S_{exe})$ is the expected markout, where S_{exe} and S_{noexe} are the terminal security prices conditional on no-execution and on execution of the limit order, respectively.
A research group at Euronext has recently published a paper on price markouts across different market venues (see references for details).

Latency

- Pricing: an asymmetric perspective.
- Trading telecommunication networks (fiber optics, microwaves).
- An example of commercial network.
- How to detect it (lagged auto correlation). Some charts.

$$C = f_{BS}(S, \sigma, K, T, r, \ldots)$$

$$F = Se^{rt} - D$$

Recap

- Market Micro-structure (COB, Roll's model)
- Pricing (APT, Fundamental, Statistical/Data Driven/PCA)
- Risk Management (Inventory control)
- Optimal Execution (TWAP, VWAP, Markouts)
- Automation (Latency problem, Filters, Regulation)