Budapest University of Technology and Economics

Faculty of Natural Sciences

Pipeline for data-driven sequential investing decisions based on theoretical optimums of trading on stochastic processes

Master's Thesis

Supervisors:
Miklós Rásonyi

Author:
Máté Matúz

Budapest, 2023

Contents

1 Introduction

We will try to find exploitable trading opportunities in commodity markets setting up the problem as the following. An agent wants to find the best achievable investment decision with the following set of actions:

1. Select an economic model for the decision-making criteria.
2. Find a suitable parametric model for the asset dynamics incorporating expert knowledge.
3. Find some theoretical optimum guarantees for the chosen dynamics and criteria. And choose a null model that we will race against.
4. Generalize the theoretical model further by adding hyperparameters if it is not ready to use in an empirical setting.
5. Fix a datastream and train the model, repeat 4. If results are unsatisfactory and choose another set of hyperparameters.
6. Fix the decision on the test period and evaluate the results. Compare with the null model.

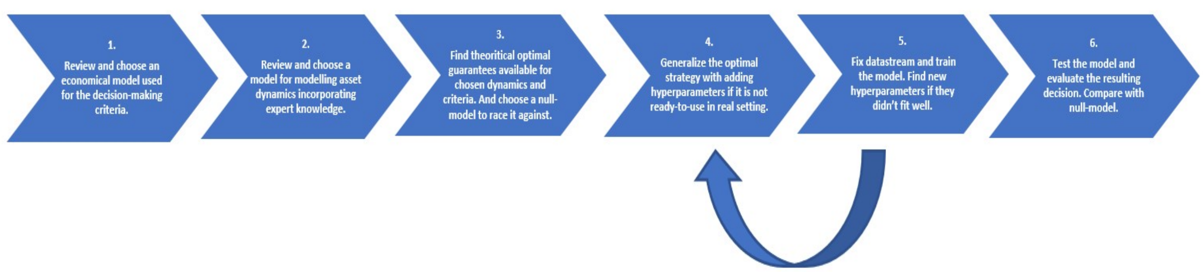

Figure 1: The decision-making pipeline

This is a rather natural idea. In the face of an economic decision with uncertain payoffs, meaning when one doesn't know the underlying probability distribution, the first step might

be for an agent to review the theories and data already available. Then try to construct a well-performing strategy starting from some idealistic setup with optimality guarantees as a good basis. Then refine it to suit a real-world setup. And finally infer some probability distributions/statistics based on the data to reduce the uncertainty of the problem.
The decisions of the agent will be made in 2 layers of time granularity.

1. In the first phase, the agent according to a model executes daily actions in a fixed window of 1 year.

2. In the second phase we look at this problem from a 1-year perspective where decisions are among the models that dictate the daily actions.

In an ever-changing market environment, we have to limit the time scope of our decisions, so we are able to evaluate and refine the decisions. There is a trade-off, ... A heuristic choice of a 1-year window worked well.
In the spirit of the new age of data-driven decision-making software, we will develop our solution in Python and try to automatize it as much as possible such that it will be in an "almost deployable" state. Meaning it is enough to provide it with data and it can retrieve a decision for the next year automatically.

2 Exploration of available theories

2.1 The economical model

We will choose from the following 3 models. These already examined the question of an economic agent's optimal decision criteria and came up with different answers. We will make comparisons and argue for and against them:

1. Neoclassical economics

2. Recursive economics

3. Ergodic economics

...First, let's introduce utility functions...

Example 2.1. St. Petersburg Paradox: We want to participate in a game of flipping coins with rules:

1. Our initial stake is 2. If we flip a head we must double our stake and keep flipping until seeing a tail.

2. If we flipped tail the game stops and we win double our stake.

Now we can see that the expected value of this game is:

$$\mathbb{E} = \frac{1}{2}2 + \frac{1}{4}4 + \cdots = \infty \tag{1}$$

... Another solution later used the logarithm of the payoff U achieved by playing the game. Namely,

$$\mathbb{E}\log U = \sum_{k=1}^{\infty} \frac{1}{2^k} \log 2^k = \sum_{k=1}^{\infty} \frac{k}{2^k} \log 2 < \infty \tag{2}$$

which is summable now. Using this one can retrieve the fair price c starting from initial wealth by using that the expected incremental utility now converges to a finite value.

$$\Delta\mathbb{E}\log U = \sum_{k=1}^{\infty} \frac{1}{2^k} [\log(w - c + 2^k) - \log w] < \infty \tag{3}$$

This means we can find such a c that this quantity changes sign so any root finding algorithm can tell us the solution. Two things we must observe here:

1. The fair price depends on the initial starting wealth under the logarithmic utility function. We will return to this point later.

2. The fair price grows fairly slowly with the starting wealth. At 1,000,000 one will be willing to pay up to 20.88, and a person starting with 1,000 will pay up to 10.95.

These investigations in the expected utility hypothesis together with rational agent theory served as the breeding ground for neoclassical economics. I will make sure the agent's action satisfies the following dependencies of the action taken by the rational agent. [8]

1. **Preferences of the agent.** We will do this by actually quantifying them with the help of a utility function.

2. **The agent's information of its environment including any available past information.** We achieve this through our research before making the decision.

3. **The set of actions available.** We will use a simple set of 2 actions. Namely, selling and buying. In future work, I want to explore a more general action space which includes holding and no-action. For example, the no-action criteria in a mean-reverting trading strategy could be described by putting a confidence band on the unobservable mean the asset reverts to and then making actions only if we have some amount of confidence say at the 50-th percentile level that a deviation from the mean actually happened.

4. **The estimated or actual benefits and the chances of success of the actions.** This part will be achieved by parametrically modeling the dynamics of our asset. But actually, this is not a necessity. In future work, I would also like to explore generalizing this step by using the very exciting field of reinforcement learning [12] where agents using, for example, the theory of Q-learning can learn the value of actions without making any kind of assumptions on distributions simply by observing and interacting with the environment.

2.1.1 Neoclassical economic model

E. Roy Weintraub. gave the following 3 assumptions characterizing the neoclassical approach: [14]

1. An agent has rational preferences between outcomes that can be identified and associated with values.

2. The agent maximizes utility.

3. The agents act independently on the basis of full and relevant information.

Of course, these are idealistic assumptions but served as a good base for the dominating economic theory together with Keynesian economics throughout 1950-70.
We will focus on the first and the second assumption. We assume the first assumption to hold in all the reviewed economic models but as we mentioned it can also be circumvented. In the neoclassical approach 2. refers to maximizing the 1-period expected utility. Meaning if we were in the setting of making daily trading decisions this would correspond to some kind of "greedy" decision algorithm. Later we will use this criterion when summarizing the 1-year results of our approach in the train phase to make an ultimate decision with the model that we act upon in the test period.
Namely, find the decision for period T

$$\Phi_T^* \quad : \quad V(\Phi_T^*) = \sup_{\Phi_T \in \Phi} V(\Phi_T) \tag{4}$$

where V is our preferred value function that aids the agent in selecting the decision for the test period.
For the agent to be able to maximize utility we will need to fix one. We choose from the following 2 families:

1. Constant absolute risk aversion family (CARA)
2. Constant relative risk aversion family (CRRA)

The utility functions in CARA also called exponential utility functions assume the form:

$$U(c) = \begin{cases} (1 - e^{-ac})/a & a \neq 0 \\ c & a = 0 \end{cases} \tag{5}$$

Where $a \geq 0$ is the risk aversion parameter and c is just a variable we control directly or indirectly. As we see in the risk-neutral case of $a = 0$ we arrive at a possibly unbounded problem. We note that excluding the constant term 1 and then maximizing the expected utility gives the same result as maximizing the expected utility of

$$U(c) = -e^{-ac}/a \tag{6}$$

Later we will just use this simpler formula.
The functions in CRRA also called isoelastic/power utility functions assume the form:

$$U(c) = \begin{cases} \frac{c^{1-\eta}}{1-\eta} & \eta \geq 0, \eta \neq 1 \\ \ln(c) & \eta = 1 \end{cases} \tag{7}$$

Where $\eta \geq 0$ is the risk aversion parameter. We have already seen the $\eta = 0$ case used in Example 2.1. The main difference between the two families will be that in CARA in an investing decision our initial levels of wealth won't affect the control variable which we will demonstrate in an example in Section 2.2. On the contrary, we have seen in example 2.1 in the CRRA case the initial wealth actually affected the solution, i.e. the price paid scaled with the initial wealth.
The neoclassical approach in the daily trading case would give the simple trading decision where the expectation is taken under the chosen dynamics $\mathbb{P}$ of the tradeable asset $X_t : \Omega \longrightarrow \mathbb{R}$, $t \in \mathbb{N}$ and ϕ_t is the policy i.e. the amount to buy or sell on day t with an underlying filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{P})$, $X_t \in \mathcal{F}_t$ measurable and $\phi_t \in \mathcal{F}_{t-1}$.

$$\phi_t^* = \arg\sup_{\phi_t \in \mathbb{R}} \mathbb{E}U(\Delta W_t^{\phi_t} | \mathcal{F}_t) \tag{8}$$

where $\Delta W_t^{\phi_t}$ is the amount gained by choosing action ϕ_t in period $(t, t+1)$
We will present a nice example using the exponential utility that shows us the tractability gained in calculation using this function.

Example 2.2. Multi-asset portfolio
We want to optimize the portfolio allocation by maximizing $\mathbb{E}[-e^{-aW}]$ the exponential utility of the final wealth W subject to

$$W = x^T r + (W_0 - x^T k) \cdot r_f \tag{9}$$

...

$$x^* = \frac{1}{a} V^{-1} (\mu - r_f \cdot k). \tag{10}$$

Here we can observe the contrast compared to the CRRA family. The optimal holding of the risky asset x^* doesn't depend on our current wealth. Meaning if were to be able to use this strategy successfully growing our wealth in consecutive periods. The share of the risky asset in our portfolio would just keep getting smaller and smaller. Therefore the **"exponential agent"** is a very careful one. Due to the tractability of closed-form results, it can have an advantage over the CRRA family.

Another famous example whose roots can be found in the neoclassical approach is the mean-variance portfolio problem also called Modern Portfolio theory.

Example 2.3. Markowitz-portfolio
Economist Harry Markowitz introduced the theory in 1952. ... [5].

$$w^T \Sigma w - q \times R^T w \tag{11}$$

where w is a vector of portfolio weights and $\sum_i w_i = 1$. The weights can be negative. Σ is the covariance matrix for the returns on the assets in the portfolio. $q \geq 0$ is a "risk tolerance" factor, where 0 results in the portfolio with minimal risk and ∞ results in the portfolio infinitely far out on the frontier with both expected return and risk unbounded. R is a vector of expected returns. $w^T \Sigma w$ is the variance of portfolio return. $R^T w$ is the expected return on the portfolio.
...

We note that this theory is not suitable to answer finer questions such as, what is the optimal series of actions in consecutive terms?

2.1.2 Recursive economics

...
Recursive economics can be used to answer questions such as, what is the optimal series of actions to maximize utility in a multi-period setting? Meaning to find $\phi_t^* : \mathbb{N} \to \mathbb{R}$ policy function such that

$$\phi_t^* = \underset{\phi_t \in \Phi}{\arg\sup}\, \mathbb{E}[U(W_T^{\phi_t}] \tag{12}$$

Where Φ is the set of all possible policy functions. And $W_T^{\phi_t}$ is the terminal wealth achieved by the series of actions ϕ_t taken in periods $[1, \ldots, T]$.

...

Example 2.4. Merton's portfolio problem [7]
This is a problem in continuous time finance. The agent has a lifetime T. Three actions are available at each time t investing π_t fraction in the risky asset, 1-π_t in the risk-free asset and consuming c_t amount. The agent's goal is to invest and consume in such a way, that it leaves a sufficient amount of wealth behind also called bequest W_T, and maximizes its utility from consumption:

$$\max \mathbb{E}\left[\int_0^T e^{-\rho s} u(c_s)\, ds + \epsilon^\eta e^{-\rho T} u(W_T)\right] \tag{13}$$

where u is from the CRRA family with risk aversion eta. Bigger eta translates into an increased reluctance of holding stocks. ε parameterizes the desired level of bequest with 0 being no amount. ρ is the subjective discount rate.
The wealth W_t evolves according to the SDE:

$$dW_t = [(r + \pi_t(\mu - r))W_t - c_t]\, dt + W_t \pi_t \sigma\, dB_t \tag{14}$$

where r is the risk-free rate, $\mu and \sigma$ are the expected return and volatility of the stock market and dB_t is the increment of the Wiener process.
We want to show an idea behind the solution which breaks down the problem into a series of discrete-time problems according to the dynamical programming approach...:

$$\max \mathbb{E}_0 \left\{\sum_{t=0}^{T-dt} \int_t^{t+dt} e^{-\rho s} u(c_s)\, ds + \epsilon^\eta e^{-\rho T} u(W_T)\right\} \tag{15}$$

This can be solved after some transformations including with the Bellman equation with the terminal condition being a function of the bequest term. Giving us the solution:

$$\pi(W, t) = \frac{\mu - r}{\sigma^2 \eta}. \tag{16}$$

As we see higher risk-aversion η corresponds to lower levels of risky assets held. The fraction stays constant in variables time and wealth. ...

2.1.3 Ergodicity economics

As we know the branch of mathematics called ergodic theory explores the relationship between time averages and expected values in dynamical systems. Not surprisingly this economic model takes this question to challenge the form of expectation used in the expected utility hypothesis. Arguing that the results derived from the latter might not correspond to the time-average solution. ...
Decisions in this branch are made in the following way: [9]
Agents are assumed to maximize the time-average growth rate of their wealth $W(t)$ The functional form of the growth rate, g, depends on the wealth process $W(t)$. In general, a growth rate takes the form

$$g = \frac{\Delta v(W)}{\Delta t} \tag{17}$$

,
where the function $v(x)$ linearizes $W(t)$, such that growth rates evaluated at different times can be compared. While growth processes $W(t)$ generally violate ergodicity (e.g. by having a positive chance of ruin), their growth rates can still be ergodic. [9] gives the conditions for this ergodicity to hold. In this case, the time-average growth rate, g_t can be computed as the rate of change of the expected value $v(x)$, i.e.

$$g_t = \frac{E[\Delta v(W)]}{\Delta t} \tag{18}$$

The connection to the expected utility hypothesis is that for a given wealth dynamics after finding $v(W)$ we can identify this function as the utility function $u(W)$ we usually maximize in the previous economic models.
How do we derive decisions from this theory?

Example 2.5. Merton's problem from ergodic perspective We will now pose a different question. The agent now wants to maximize the time-average growth instead of any specific utility. And we only know the dynamics of our portfolio.

In example 2.4 examining Merton's problem we implicitly assumed the risky asset follows a Geometric Brownian Motion with drift rate μ and diffusion rate σ. If we forego the option to consume in the example and simplify to a constant portfolio in time weight it gives us the wealth dynamics

$$dW_t = [(r + \pi(\mu - r))W_t]\, dt + W_t \pi\sigma\, dB_t \tag{19}$$

that is again a GBM. Renaming the terms gives us

$$dW_t = \tilde{\mu} W_t\, dt + W_t \tilde{\sigma}\, dB_t \tag{20}$$

where $\tilde{\mu} = r + \pi(\mu - r)$ and $\tilde{\sigma} = \pi\sigma$
The solution is well known

$$W(t) = W(0) \exp\left(\left(\tilde{\mu} - \frac{\tilde{\sigma}^2}{2}\right) t + \tilde{\sigma} W_t\right) \tag{21}$$

The linearization $v(x)$ in this case is $v(x) = \ln(x)$ giving us the time average growth-rate

$$g_t = \frac{E[\Delta v(x)]}{\Delta t} = \tilde{\mu} - \frac{\tilde{\sigma}^2}{2} \tag{22}$$

It follows that for geometric Brownian motion, maximizing the rate of change in the logarithmic utility function, $u(x) = \ln(x)$, is equivalent to maximizing the time-average growth rate of wealth. This actually fixes

$$\eta = 1 \tag{23}$$

in (20). Meaning there is one utility function that we can maximize to solve this new question. Giving us the time-average growth optimal risky asset share of

$$\pi = \frac{\mu - r}{\sigma^2}. \tag{24}$$

.

... We know can actually calculate which utility functions we need to maximize for specific objectives.
... resembling the Kelly Criterion [4].

2.2 Review of models of asset dynamics and corresponding optimality guarantees

Our aim here is to non-exhaustively explore the available choices and choose one that fits our criteria the most.
We can cover all possible dynamics with two sets:

1. Stationary or non-stationary processes

2. Discrete time or continuous time processes

...

...

2.2.1 Discrete time, stationary processes

Maximal growth rate Using the objective of maximal average growth rate which covers decision model 2.1.3. we can look toward the works of László Györfi. [3] Together with other authors he provided the following results:

1. For the above objective maximizing the logarithmic utility is the right choice asymptotically. ...

2. Both non-parametric and parametric models of the asset price dynamics are covered. In the parametric case... In the non-parametric case, we assume only that the returns of the assets are stationary and ergodic. Then the theory gives us methods to estimate the next-period returns based on the realizations already observed such that the famous machine learning algorithm K-Nearest-Neighbour estimator and kernel-based methods. The theory also guarantees that using these estimates to maximize the logarithmic utility of the portfolio dominates the growth rate of any other portfolio almost surely in the whole family. Referred to as strong universal consistency...

During my university studies, I have explored the performance of this theory on a global set of ETF-s in the framework of the course Individual Project. It performed well and I observed that this even held in seemingly non-stationary markets. A weak argument for why might

this hold comes from the ergodicity economics results where we have seen that even on a non-stationary market with GBM dynamics maximizing the logarithmic utility still achieves the maximal growth rate. This is a seed of suspicion that this approach might even hold in larger non-stationary families as well as explaining the good empirical performance.
With this, we covered all the dynamics for the decision criteria of maximal growth rate.
Now we will explore dynamics individually and the corresponding theoretical guarantees.

Autoregressive processes Stationary autoregressive process:

$$X_t = \phi X_{t-1} + \sigma\varepsilon_t, \quad \phi \in (-1,1), \sigma \in \mathbb{R}^+ \tag{25}$$

With ε_t being a standard normal $\mathcal{N}(0,1)$
In this case, under the recursive economics decision criteria, we have finite-sample optimal guarantees under the exponential utility function. Namely in [1] we get an explicit recipe to construct a strategy using past information and maximize the exponential utility in finite time. The main draw is that it allows for short selling and buying also and guarantees finite-sample optimums.

Log-Markov process In [6] the following model is introduced:

$$S_t = \exp(X_t), \quad X_t - X_{t-1} = \mu(X_{t-1}) + \sigma(X_{t-1})\varepsilon_t,\ quadt \geq 1 \tag{26}$$

...

2.2.2 Discrete time, non-stationary processes

In this case, we can find many models. Stochastic trend models also called integrating processes. Trend stationary models, where we assume a linear trend and an added noise. Ultimately these models suggest the possibility of unbounded growth. ...

Hidden Markov Model ... In future work, I want to explore the usage of this model in mean-reverting strategies. [11] ...The theory can be used efficiently in a high-frequency trading setting as well.

2.2.3 Continous time, stationary processes

We explore the famous dynamics called the Ornstein-Uhlenbeck process. Using the stochastic differential equation notation it can be written as

$$dx_t = \theta(\mu - x_t)dt + \sigma dW_t, \quad \mu, \theta, \sigma \in \mathbb{R} \tag{27}$$

where dW_t denotes the increment of the usual Wiener process.
Results:
In [13] the authors solve the problem of optimal timing in markets analytically in a wide area of parametric models. Amongst them they give analytical solutions on trading with OU process, how to find cointegrating pairs with loglikelihood methods, also exit strategies, and entry strategies.
In [2] we can see the problem of constructing mean reverting portfolios from another perspective with the emphasis placed on trying to quantify and then maximize mean reversion using a VAR(1) process then selecting sparse linear combinations of the components of this process which they then model as OU-processes. ... In conclusion, they translated finding predictable portfolios into an eigenvalue problem.
...

2.2.4 Continous time, non-stationary processes

We have already seen some results on the most famous non-stationary dynamics GBM.
...We have seen 2 analytical results...
Still, the possibility of unbounded growth exists so ultimately we decided against this category of models.

2.3 Choice

Ultimately we have chosen the "exponential agent" model with autoregressive process as a dynamics that gives theoretical guarantees in a finite-time horizon and also allows for short selling and buying. We will later extend the model to hold the form

$$U_t = c + X_t, \tag{28}$$

where X_t is just the previous autoregressive process and $c \in \mathbb{R}$ a constant which we fluctuate around.

2.4 Null model

The null model will be called the Equivalent AR1 model. It assumes the same dynamics:

$$U_t = c + X_t, X_t = \gamma X_{t-1} + \varepsilon_t \tag{29}$$

The null model will use the same amount of ”resources” on a 1-year time interval, but it misses information the other model takes into account. Namely the magnitude of the deviations. If we think about the optimal strategy in terms of optimally distributing resources along a time series of returns (which is basically the same as loading and offloading positions on 1-day intervals). Then the natural comparison is a model which can't use the information on how to optimally distribute those resources. Still an AR(1) model can give predictions using conditional expectation on the predicted sign of the next day's return.
We will now formalize this.

Definition 2.1. *Cost Series*
The cost series $\{C_t\}$ *of asset* $\{U_t\}, U_t = c + X_t$ *using strategy* $\phi_t^T(z)$ *is*

$$C_{t-1} = \phi_t^T(X_{t-1}) \cdot U_{t-1} \tag{30}$$

Definition 2.2. *Cost Of Business (COB)*
COB is the sum of the absolute value of all the positions the optimal strategy held

$$COB(\{C_t\}) = \sum_{t=1}^{T} |C_t| \tag{31}$$

Definition 2.3. *Equivalent average position resource $\bar{C}$ is*

$$\bar{C} = COB(\{C_t\}_{t=1}^{T})/T \tag{32}$$

Example 2.6. Example of prediction and position size of Equivalent AR1

$$\mathbb{E}(U_t|U_{t-1}) = c + \gamma X_{t-1} \tag{33}$$

where c is now assumed to be known. This means that on the event $U_{t_1} > c$ we must sell otherwise we buy.
This means for the equivalent strategy $\phi_t^{T,AR1}$

$$\phi_t^{T,AR1}(U_{t-1}) = -\frac{U_{t-1}}{\bar{C}}\mathbb{1}\{U_{t-1} > c\} + \frac{U_{t-1}}{\bar{C}}\mathbb{1}\{U_{t-1} < c\} \tag{34}$$

where $\mathbb{1}(\cdot)$ is the usual indicator function.

2.5 Extension and improvement

Now that we fixed the dynamics of the asset X_t:

$$X_t = \gamma X_{t-1} + \sigma\varepsilon_t \tag{35}$$

And the T-optimal strategy corresponding to it for the daily decision granularity in the 1-year periods $t \in \{1, \dots, T\}$:

$$\phi_t^T(z) = \frac{(\gamma - 1)z}{\sigma^2}\theta_t^T \tag{36}$$

where $\theta_t^T = 1 - (T - t)(\gamma - 1)$.
We can start to extend the model with hyperparameters. First a corollary of **Theorem 2.1** in [1]

Corollary 1. *If $U_t = c + X_t$ is the observable process, then the optimal strategy for U_t which is the sum of an $AR(1)$ noise and a mean $c \in \mathcal{R}$.*

$$\phi_t^T(X_{t-1}) = \frac{(\gamma-1)X_{t-1}}{\sigma^2}\theta_t^T \tag{37}$$

where $X_t = U_t - c$*.*

Proof. Starting from (23) we just use $X_t = U_t - c$ and show that nothing changes when minimizing the exponent. We are in the first step $T = 1$ case. We are now trading on the observable process U_t, $\beta = \gamma - 1$.

$$L_1^\phi = \phi_1(U_1 - U_0) = \phi_1(X_1 - X_0) = \phi_1(\beta X_0 + \sigma * \varepsilon_1) \tag{38}$$

We got back to the original case. Now we only have to take care of the filtrations used for conditioning in (24).

$$\mathbb{E}[e^{-\phi_1(\beta X_0+\sigma\varepsilon_1)} \mid U_0 = z] = \mathbb{E}[e^{-\phi_1(\beta X_0+\sigma\varepsilon_1)} \mid Z_0 = z - c] = \mathbb{E}[e^{-\phi_1(\beta X_0+\sigma\varepsilon_1)} \mid Z_0 = \tilde{z}] \tag{39}$$

Now the rest of the proof holds as before for this new $\tilde{z} \in \mathbb{R}$ constant. And at the end, we substitute $\tilde{z} = z - c$

□

2.5.1 First iteration

...

2.5.2 Second iteration

We will adapt to the dynamic environment of the markets. We use rolling mean to calculate a theoretical mean that we revert to. We will use an expanding window on the calculation of variance and also on the γ coefficient.

Definition 2.4. *Rolling mean process*

The rolling mean process c_t^h of X_t will be defined as

$$c_t^h = \frac{1}{h} \sum_{i=t-h}^{t-1} X_i, t \in \{30, \ldots, T\} \tag{40}$$

where the first 30 days are reserved of the training period for parameters.

Definition 2.5. *Expanding variance, coefficient process σ_t^2, γ_t*
Regressing$\{X_k\}_{k=2}^{T+t}$ on $\{X_k\}_{k=1}^{T-1+t}$ gives σ_t^2, γ_t at each t.

The new dynamic is now:

$$U_t = X_t + c_t^h, \quad X_t = \gamma_t X_{t-1} + \sigma_t \varepsilon. \tag{41}$$

Our autocorrelation plot is now "improved".

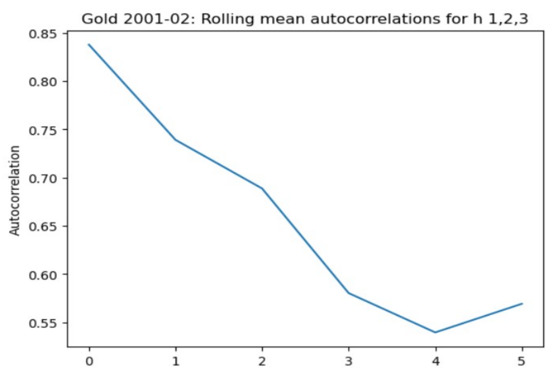

Figure 2: Gold autocorrelations

Here improvement means that directions of the predicted movements now more confidently match up with realized movements.

2.5.3 Is there a point to optimizing the window h?

Shortly, empirically yes.

...

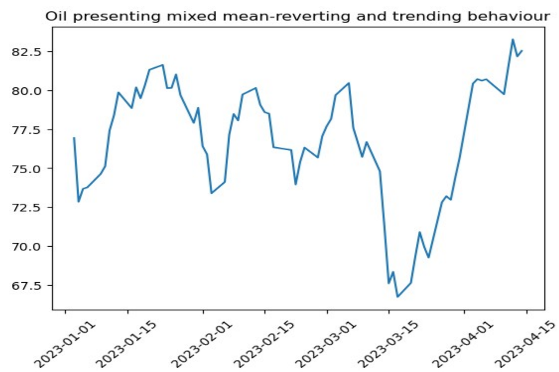

Figure 3: Mixed behavior of oil

Ideally, we would want the nice periodic movement to persist. Still, we can make use of this nice separation of oscillating and trending parts the following way.
...A crude approximation...

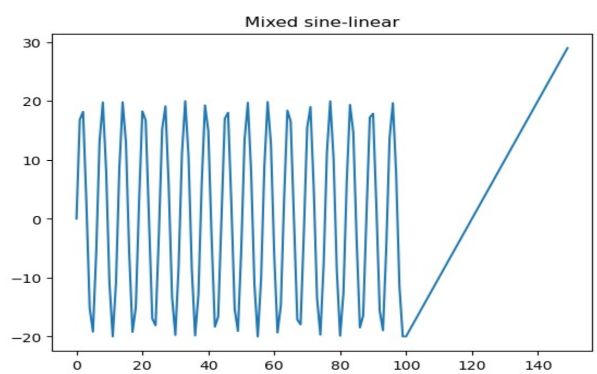

Figure 4: Sine-linear mixed

Now let's see the wealth achieved W_T as a function of window h trading on only the sine part.

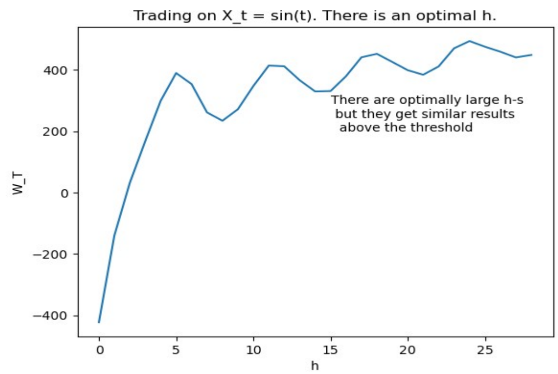

Figure 5: W_T as a function of h

As we see after a threshold the improvement takes on a slower slope. Of course, the results improve because we are less and less sensitive to the periodicity of the sine-wave, the mean predicted stays close to zero.

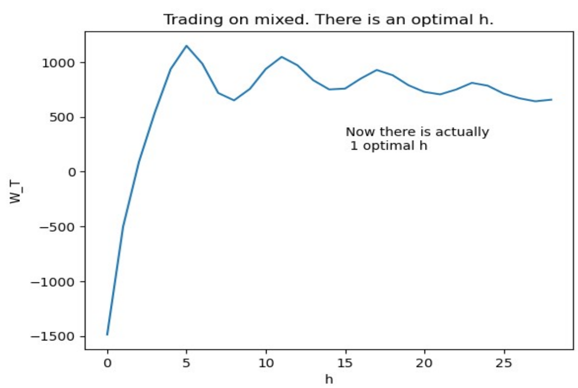

Figure 6: W_T as a function of h

After adding the linear term now there is an optimal h we can look for. This might be one of the reasons we can hope for an improvement. Of course, these are very heuristic arguments.

3 Train and test

The train datastream is fixed from 2001-2015 a 15-year training period consisting of 9 assets in 1-year increments $\{U_{t,i,k}\}$ with $k \in \{1, \ldots, 15\}$

$$i \in \{Soybean, Wheat, Coffee, Cocoa, Sugar, Corn, Crudeoil, Gold, Copper\}. \tag{42}$$

We will make one out-of-sample prediction after that for the following year. Answering the question of what is the best choice of h, and best choice asset $\{U_{t,i}\}_{16}$ to make a "safe" bet where we will define the meaning of "safe" in the form of a value function.
We will as mentioned previously use the null model as a base and start with a heuristically set window $h = 10$ in the rolling mean. Then we will use some clever iterative techniques to search in the hyperparameter space for better results.

3.1 Definitions

We will define some additional metrics that will guide our exponential agent in the decision-making process

Construction of value function using Post-Modern Portfolio Theory In section 2.1.1 we introduced the 1-period decision model MPT. ...
Post-Modern Portfolio Theory aims to correct the following limitations of MPT:

1. Variance of portfolio returns is assumed to be a correct measure of risk.
2. The investment returns of all securities and portfolios can be adequately represented by a joint elliptical distribution, such as the normal distribution.

Firstly, it has been observed that investors are more averse to lower-than-expected returns rather than the converse [10]. ... To solve this the author presented the risk metric called semivariance also called downside risk.

Definition 3.1. *Semivarince d^2*

$$d(X) = \left(\mathbb{E}[(X - \mathbb{E}[X])^2 1_{\{X \leq \mathbb{E}[X]\}}]\right)^{\frac{1}{2}} \tag{43}$$

Where X is a random variable. Note that the semivariance of the normal distribution having a skewness of 0 is just half of its variance simply by the additivity of integrals and symmetry. Therefore it gives the same value to the upside risk (we flip the inequality) and downside risk. Thus the aim here is to generalize these 2 measures. So the question evolves from how far away will I be from the mean usually to if I expect a negative return how far away will I be from the mean?

Secondly, symmetrical distributions like the normal do not adequately represent the sometimes skewed nature of returns. ... To admit both negative and positive skewness it was generalized into a new distribution the authors call the three-parameter lognormal distribution.

Definition 3.2. *3-parameter lognormal*
The 3-parameter lognormal is just a 2-param lognormal shifted by a constant γ

$$Y = X - \gamma, \quad \gamma \in \mathbb{R} \tag{44}$$

$X \sim Lognormal(\mu, \sigma^2)$.
The need for this is because otherwise the lognormal mean and standard deviation can't be set independently. Now the mean can be tuned as well independently. This is always positively skewed but we can take the negative of the distribution. The skewness is controlled by σ *only. Now together with the negative lognormal, this covers every case. And can be used for robust estimation of return distributions.*

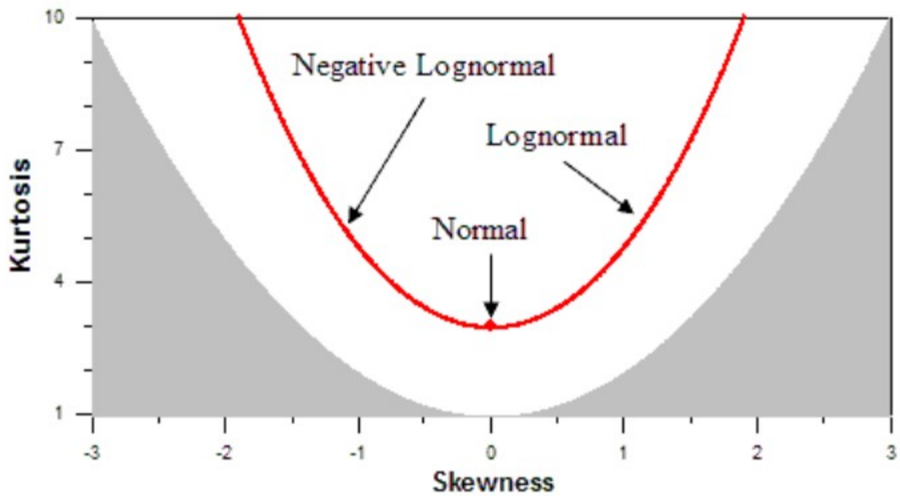

Figure 7: Skewness-Kurtosis of lognormal family

Some assumptions to note for the exponential agent

1. The exponential agent works with absolute returns. And the starting wealth will be 0.
2. This means we can take on positions by borrowing money costlessly and then paying back the borrowed amount the next day netting the price change times our position.
3. 2. is not as ideal as it sounds since if we lose at the end we have to pay back the money
4. 2. can also be viewed as no-money-down margin trading. Margin trading already exists with as high leverages as 1:500. Meaning we can take on positions holding 500 times our wealth. Of course that is a way a to quick margin call. When we open a position with 1:500 leverage we are jointly owning that position by the loan we get from the broker. Meaning our share is 1/500, a little movement in the wrong direction can take this down to 0 because every loss and every win is covered by our original share such that the borrowed money should not suffer losses. This basically makes us 500 times more sensitive to price swings. Now between those two cases 0∼0% and 1/500∼100%, there is a limit that the broker sets usually around 20-30% of the starting share of 100%. This is done in order to limit the possibility of not only losing our money but

also the broker's money. Based on this we make a metric that compares the terminal wealth to the assumed ideal "cost" of the strategy the 20% of the biggest absolute position cost. Representing the idealistic idea that somehow the losses and wins in time balance each other such that this stake is enough to cover all the positions. And we will use this as a "starting wealth" in a metric to get a percentage-wise idea of the gains also. Of course, we can't know this in advance and this is a weak point of this metric that it gives a posteriori view.

5. In a realistic setting an empirical estimate of starting wealth for an asset would be just using the historical position sizes across the years and constructing something like the well-known Value-at-Risk or VAR95 metric, so that we can cover our position sizes 95% of the cases.

6. The 20% percent assumption gives us a leverage of 1:5 which is actually pretty small. Ideally, by choosing the exponential agent we want this strategy to give consistent low-risk returns then we can scale it with margin trading and achieve massive gains with limited downside.

Definition 3.3. *Wealth process of asset* i $\{w_{t,i,j,k}\}_{t=1}^{T}, j \in \{-, ar1\}$*, in year* k

$$w_{T,i,j,k} = \sum_{t=1}^{T} \phi_t^j (U_{t-1,i,k} - c)(U_{t,i,k} - U_{t-1,i,k}), \quad w_{0,i,j,k} = 0. \tag{45}$$

Where the ϕ_t^j *is just the usual strategy in the original,* $j = -$*, and the null-model,* $j = ar1$*, cases as defined before.*

Definition 3.4. *Return process* $\Delta_{t,i,j,k}$ *of asset* i *in year* k

$$\Delta_{t,i,j,k} = \phi_t^j (U_{t-1,i,k} - c)(U_{t,i,k} - U_{t-1,i,k}) \tag{46}$$

Definition 3.5. *Average absolute return* $\hat{\Delta}_{i,j}$ *of asset* i *in year* k

$$\hat{\Delta}_{i,j,k} = \frac{1}{T} \sum_{t=1}^{T} \Delta_{t,i,j,k} \tag{47}$$

Definition 3.6. *Absolute return variance of asset i in year k*

$$\sigma^2_{i,j,k} = \frac{1}{T}\sum_{t=1}^{T}(\Delta_{t,i,j,k} - \hat{\Delta}_{i,j,k})^2 \tag{48}$$

Definition 3.7. *Value function in the 1-day granularity $V_{day,i,j,k}$ of asset i in year k is the usual Sharpe ratio*

$$V_{day,i,j,k} = \frac{\hat{\Delta}_{i,j,k}}{\sigma_{i,j,k}}\sqrt{252} \tag{49}$$

Definition 3.8. *Average yearly growth $\hat{w}_{T,i,j}$*

$$\hat{w}_{T,i,j} = \frac{1}{15}\sum_{k=1}^{15} w_{T,i,j,k} \tag{50}$$

Definition 3.9. *Yearly semivariance $d_{i,j}$ of asset i*

$$d_{i,j} = \frac{1}{n}\sum_{k=1}^{n}(w_{T,i,j,k} - \hat{w}_{T,i,j})^2 \mathbb{1}_{(w_{T,i,j,k} \leq \hat{w}_{T,i,j})} \tag{51}$$

where n is the number of lower-than-expected return years.

Definition 3.10. *Value function in the 1-year granularity $V_{year,i}$ of asset i is like the Sortino-ratio from PMPT but we don't scale it. Here we omit the j index because of the poor first-turn results of the null model we will only check the original strategy.*

$$V_{year,i} = \frac{\hat{w}_{T,i}}{d_i} \tag{52}$$

Definition 3.11. *Median yearly wealth m_i*
m_i is just the regular median taken over the yearly returns $\{w_{T,i,-,k}\}_{k=1}^{15}$ of asset i

Definition 3.12. *Biggest absolute position $M_{i,k}$ of asset i in year k*

$$M_{i,k} = \max\{|\phi_t \cdot U_{t-1,i,k}|\}_{t=1}^{T} \tag{53}$$

Definition 3.13. *20% margin returns $R_{20,i,k}$ of asset i in year k*

$$R_{20,i,k} = \frac{w_{T,i,-,k}}{M_{i,k}} \tag{54}$$

Definition 3.14. *Safer robust alternative to $V_{year,i}$ is using the median version $V_{year,i,safe}$ of asset i*

$$V_{year,i,safe} = \frac{m_i}{d_i} \tag{55}$$

3.2 Train results

We train the model as stated in 2.4.2. with a learning window of 30 days then after that continuously generate the positions and initiate trading. End results of the 1-year periods are plotted on a series of boxplots. The boxplots will contain the 5 metrics $\{R_{20,i,k}, w_{T,i,j,k}, V_{day,i,j,k}, j \in \{-, ar1\}\}$. Later we won't need to look at these just good to observe what happens on the training data once. Then we will be only interested in the test phase results which will be presented by scatter plots and tables.

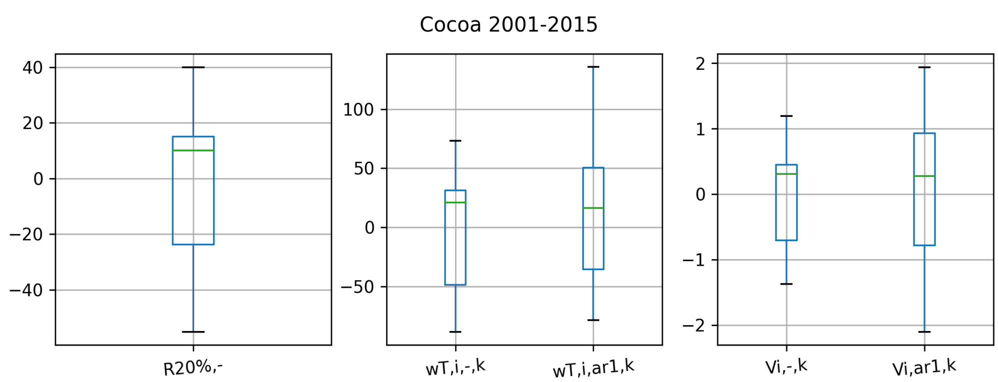

Figure 8: Cocoa

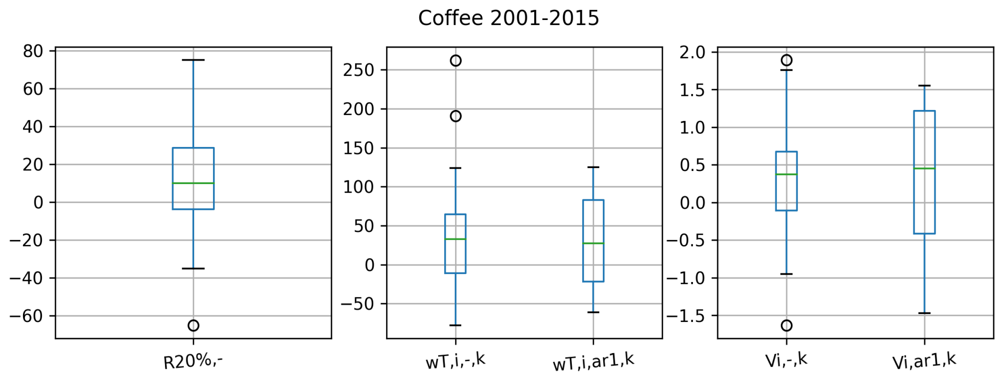

Figure 9: Coffee

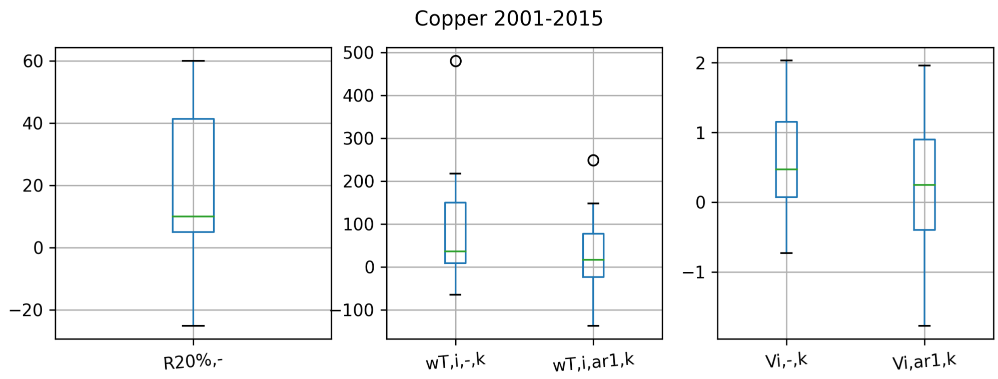

Figure 10: Copper

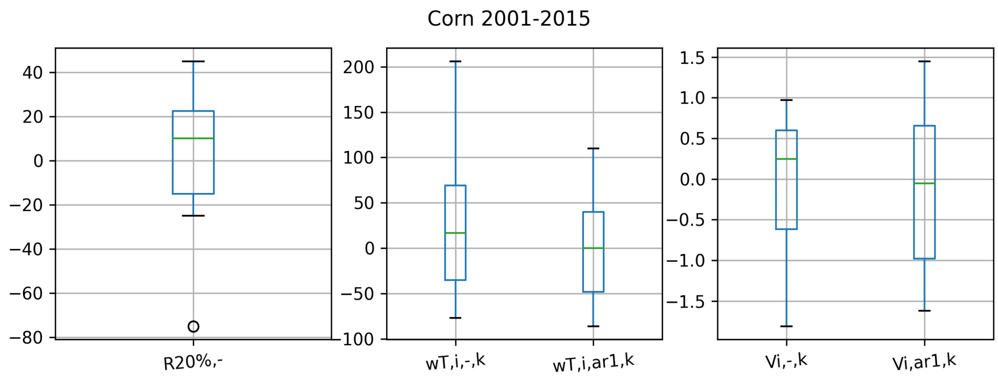

Figure 11: Corn

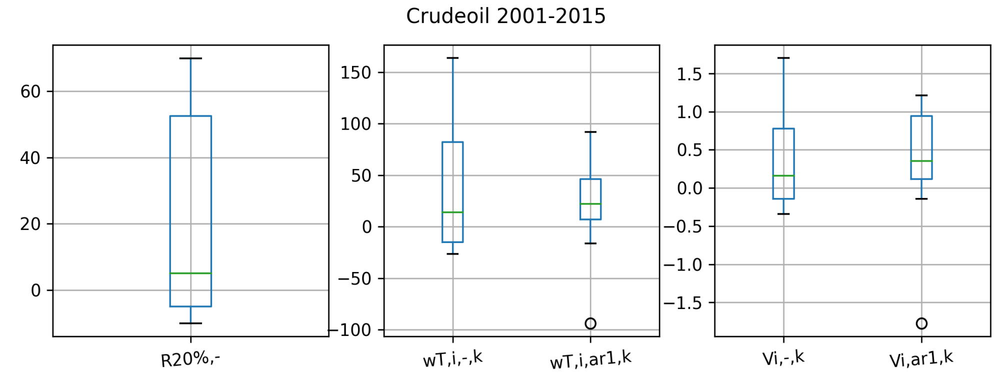

Figure 12: Crude oil

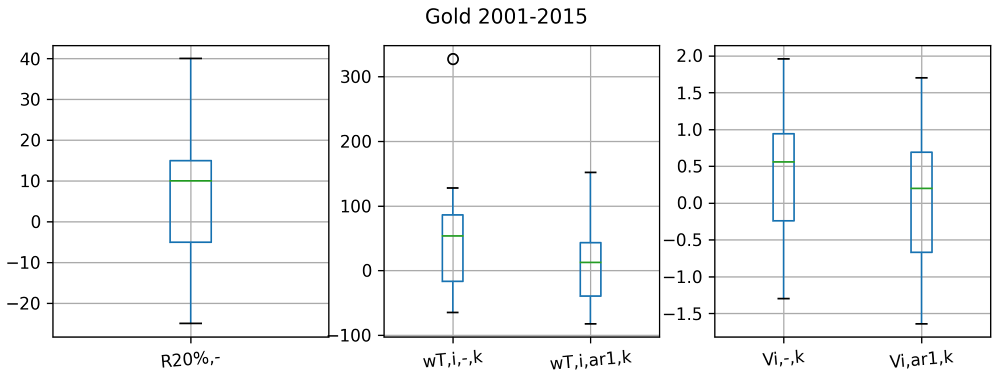

Figure 13: Gold

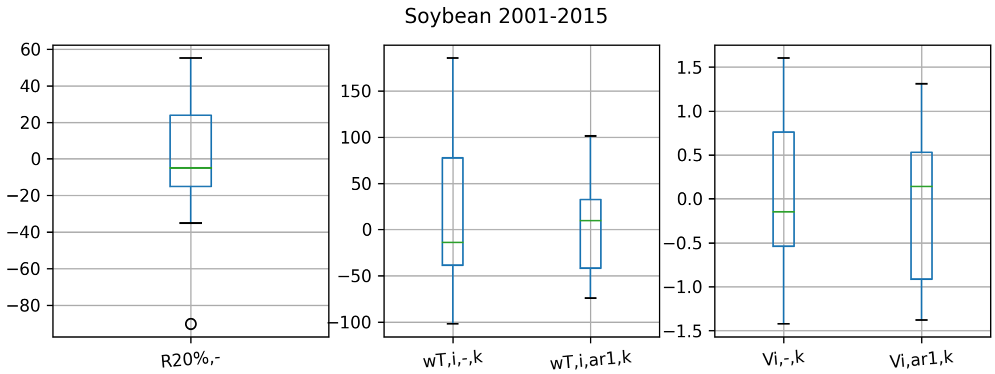

Figure 14: Cocoa

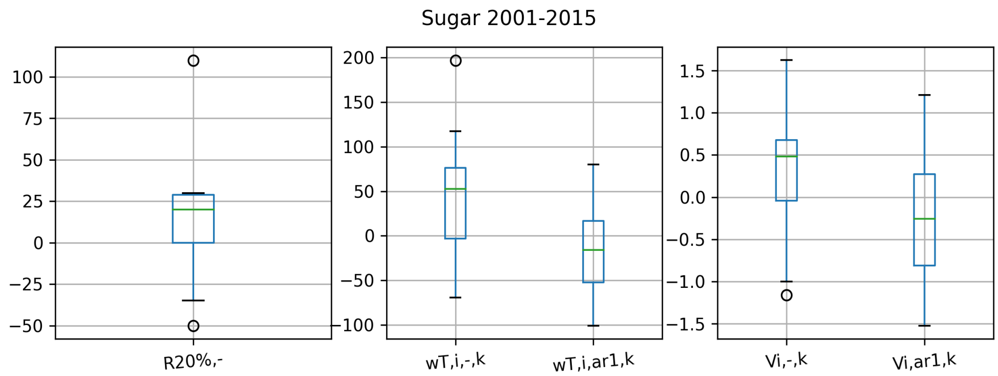

Figure 15: Cocoa

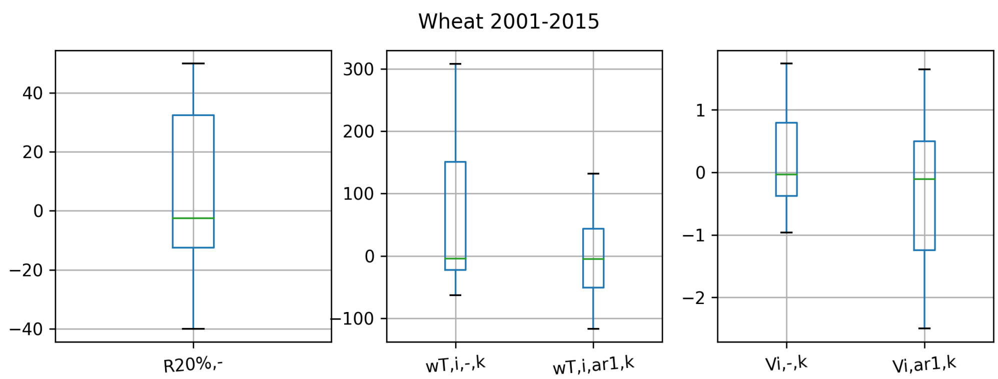

Figure 16: Cocoa

Observation of the training statistics and evaluation of the null model

1. The $R_{20,i,k}$ metric shows a positive upside with limited downside. We achieved growth as high as + 118% with the biggest loss of - 82%.

2. The null model was dominated by the exponential agent. We achieved on average bigger terminal wealth with lower semivariance and many big positive outliers not present in the null model. We make the decision as it were in hypothesis testing and say there is ”significant” empirical evidence that we have improved the performance by incorporating the information about the sizes of deviations.

3. The day-granular value functions also support the previous observations.

4. We have observed big positive outliers therefore the non-robust mean version of the value function might give distorted results.

What does our agent see? Let's see the boxplot of terminal wealths in 2 different ordering. Once according to the semivolatility-mean value $V_{year,i}$ function and once according to the safer semivolatility-median $V_{year,i,safe}$.

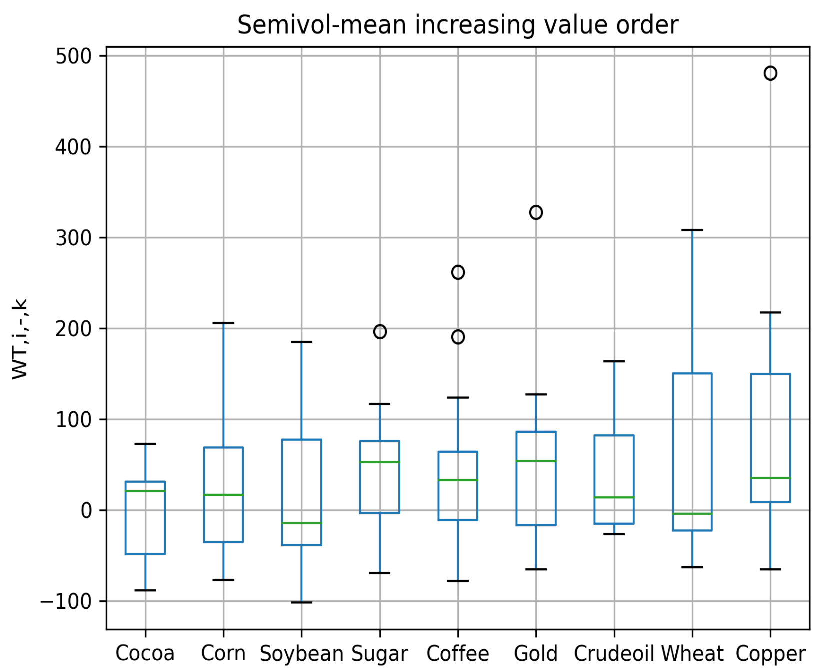

Figure 17: WT,i according to risky

We can on Figure 17 that the big positive outliers dominate the order.

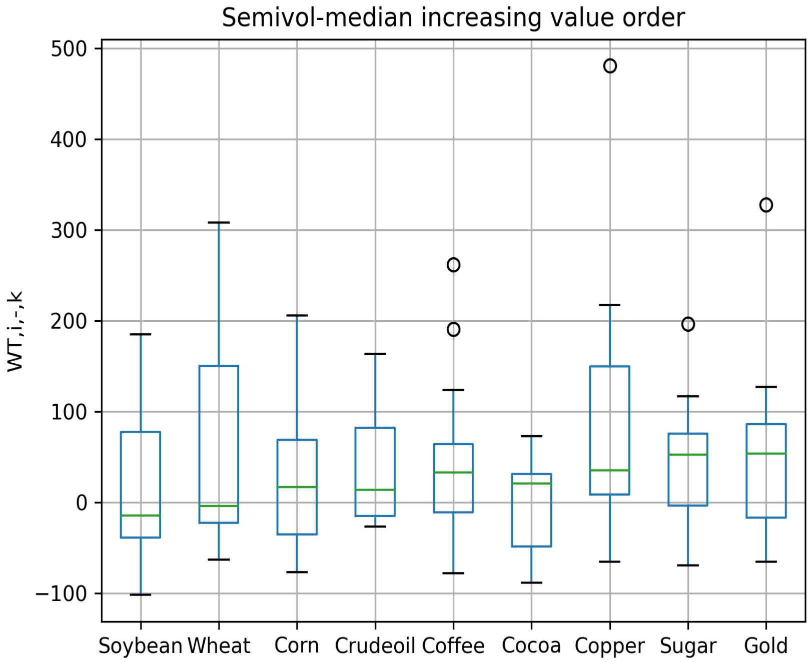

Figure 18: WT,i according to safe

The results in Figure 18 as expected are vastly different from the preference for high consistent returns replacing Copper with Gold as the best choice.

We can observe this in the actual scatterplots of returns against semivolatility.

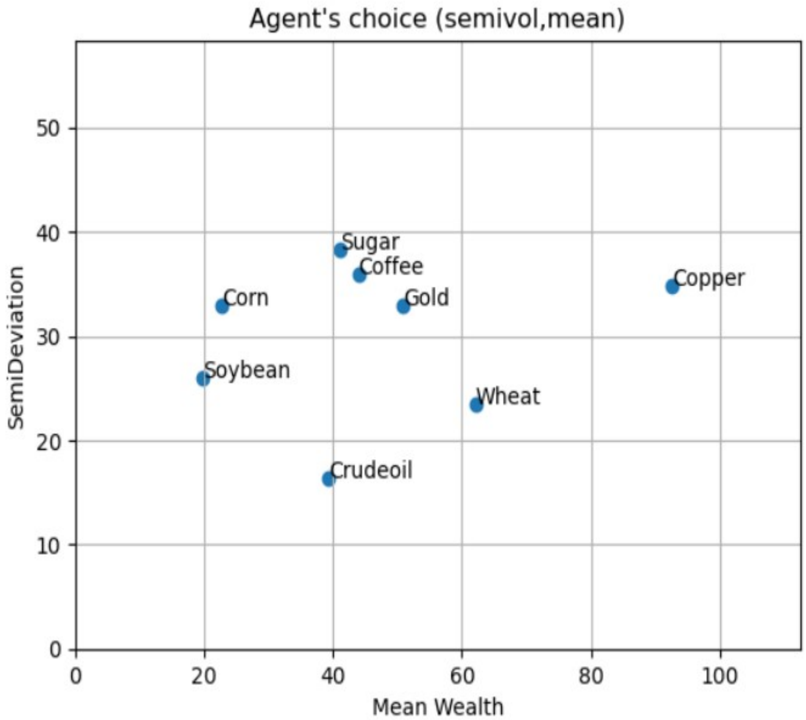

Figure 19: Agent's choice, view in the Semivol-Mean world

This is just a very expressive view. It is interesting to see the abstract world the agent views the decision through.

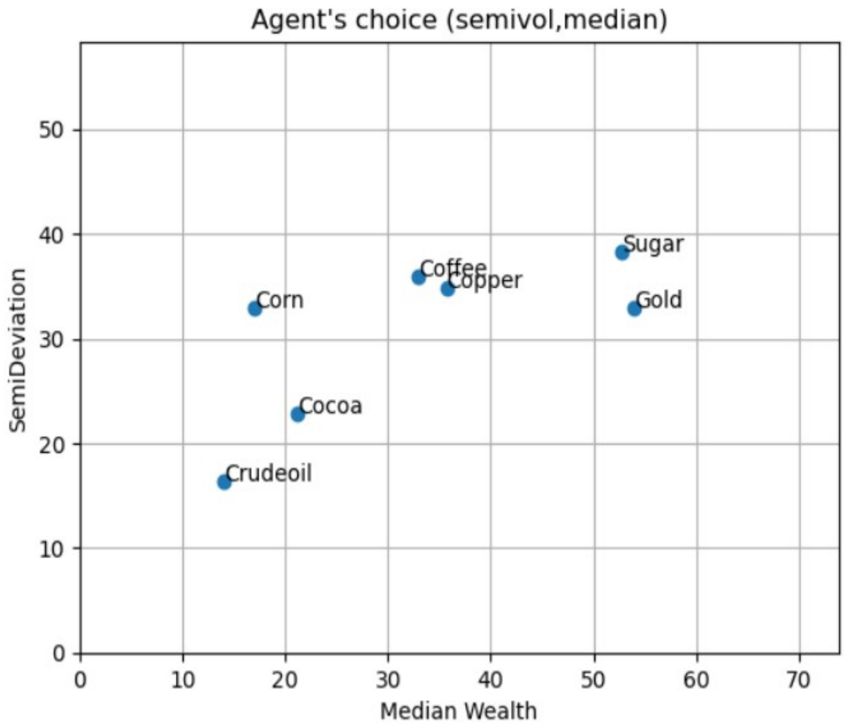

Figure 20: Agent's choice, view in the Semivol-Median world

Now we are ready to discard the null model and make out-of-sample predictions with the base model together with the 2 yearly metrics with preset window h = 10.

3.3 Out-of-sample tests

Non-optimized rolling predictions We will do rolling predictions, by rolling the 15-year training window also until 2023 starting from 2016. There won't be any point in adding up the results. We just want to observe the stability of the predictions. The results are as follows.

	Risky	Risky Asset	Safe	Safe Asset
2016	-53.27	Copper	43.47	Gold
2017	59.08	Copper	-73.94	Gold
2018	-0.18	Crudeoil	434.77	Gold
2019	250.58	Copper	15.26	Cocoa
2020	-45.25	Copper	-45.25	Copper
2021	12.15	Gold	117.05	Cocoa
2022	353.88	Crudeoil	342.76	Cocoa
2023	-12.61	Crudeoil	-14.33	Cocoa

Figure 21: Test Results

Both metrics gave very good results with big positive outliers and limited downside. The exponential agent was able to extract good opportunities and when there weren't any minimize losses. 2020 was the year of COVID so naturally, we expected losses there because there were big negative and positive trends in the commodities and we prefer stationary mean-reverting regimes still compared to the positive results they are neglectable. The interesting thing is they managed to find different big positive outliers. Somehow using both metrics can give a more stable steady gain.

Now we optimize the hyperparameter h, the rolling window The rules are:

1. Since both value functions performed well now we choose the safe value function only
2. We try to predict what is the best window for the next year by training the model on a 15-year period with h=10 initially, then applying the exponential agent's decision to get an asset to invest in next year. Then we optimize the rolling window only for that asset. Then we calculate the next training period with that new h now. And see if this results in a better guess next year. We are theorizing that some kind of interesting dynamical property might persist that we can capture and exploit.
3. Now this strategy becomes much more path dependent so the agent should choose bigger consecutive windows sticking with the strategy to evaluate it properly.

The results:

	Optimal window for the next year	Safe	Safe Asset
2016	2	43.468067	Gold
2017	2	257.480763	Copper
2018	1	336.697437	Copper
2019	1	1018.574952	Crudeoil
2020	1	40790.541242	Crudeoil
2021	1	2733.539329	Crudeoil
2022	29	-875.643683	Copper
2023	4	42.326287	Copper

Figure 22: Predicted-h-optimized results

As a sanity check, we start from the same Gold prediction. The results are quite interesting. We see that in 2017 the agent choose the same asset as the risky choice previously but was able to improve it and guess the optimal window of 2 for the next year achieving 500% of the wealth on the asset Copper. Then the agent chooses a subjectively risky strategy. It sets the rolling window to 1 therefore always measuring the deviation of day t from the day t-1. It achieves big hits on the Crude Oil asset with this strategy as big as 100 times the previous non-optimized positive outliers. Interestingly this big positive outlier happened in the COVID year, meaning the agent adapted quite well with a very short window to the turbulations of that period. We only see a downside when the other two models were able to perform in 2022 but h-optimized achieved a relatively bigger loss pursuing the risky small window strategy. It is amply covered though by the wins. Overall we had fewer negative return years only 1. Again we achieved small downside risk with very big positive outliers. The strategy looks ripened to be scaled with margin trading and we are motivated to upgrade the model to fit into a more realistic setting that considers starting wealth also.

This was a novel idea for prediction by attempting to capture some persisting dynamical properties through the size of the window parameter. Since we are constrained by the small number of data points available we can only say that there is positive evidence for upgrading and exploring these models. The consistency and stability even in such a small sample is remarkable. The results give a reason for some modest satisfaction/motivation.

4 Summary

We achieved satisfactory results and have seen that the value functions constructed helped the agent make optimal decisions with a statistical edge toward positive results.

Future work Many questions are open still. The composition of the thesis gives a recipe to make an autonomous agent to make optimal investing decisions. The main theme was incorporating information cleverly to gain a statistical edge which we have achieved.
Further generalizations:

1. As we mentioned the action space can be made larger with 4 actions instead of 2 incorporating holding, and no-action.
2. We might try to construct and analyze distributions from the year-end results further.
3. We can try to prove the result in continuous time.
4. We can try to set up the exponential optimum problem as a stochastic programming one and try to solve it in a multi-asset setting where the control variable is again the position weights but we constrain it to sum to 1. Then we can achieve percentage-wise results.
5. Reinforcement learning is an open avenue and instead of solving the problem parametrically, we can let an agent learn the distributions and payoffs of corresponding actions in that space.
6. An interesting empirical question is whether the rolling mean could be outperformed by the use of the rolling median.
7. The extension of the analytical result to VAR1 is an open question.
8. We could find finite-sample maximal growth rate results.
9. We could examine how far away are we from optimality using the rolling mean instead of the constant one.

10. We can race it against other null models like the multi-asset example.
11. We can use diversification to achieve lower semivariance with more stable returns. This could turn out to be a computationally complex search examining mixed strategies in the training set.

5 Closing and thanks

I would like to thank all my teachers for the honest effort and attention they put into answering my many questions and for their guidance throughout my studies. I would like to thank my supervisor for the opportunity to work together.

References

[1] S. Deák and M. Rásonyi. An explicit solution for optimal investment problems with autoregressive prices and exponential utility. 2015.

[2] N. Fogarasi and J. Levendovszky. Sparse, mean reverting portfolio selection using simulated annealing. 2013.

[3] L. Győrfi, G. Ottucsák, and A. Urbán. Empirical log-optimal portfolio selections: a survey. 2012.

[4] J. L. Kelly. A new interpretation of information rate. 1956.

[5] H. Markowitz. Portfolio selection. 1952.

[6] M. L. D. Mbele-Bidima and M. Rásonyi. Asymptotic exponential arbitrage and utility-based asymptotic arbitrage in markovian models of financial markets. *Acta Appl. Math. 138*, 2015.

[7] R. C. Merton. Lifetime portfolio selection under uncertainty: the continuous-time case. 1969.

[8] J. V. Neumann and O. Morgenstern. Theory of games and economic behaviour. *Princeton University Press: Princeton, NJ*, 1944.

[9] O. Peters and A. Adamou. The time interpretation of expected utility theory. 2021.

[10] W. F. Sharpe. Capital asset prices: A theory of market equilibrium under consideration of risk. *Journal of Finance*, 1964.

[11] R. Sipos, A. Ceffer, and J. Levendovszky. Parallel optimization of sparse portfolios with ar-hmms. 2017.

[12] R. S. Sutton and A. G. Bart. Reinforcement learning: An introduction. 2014.

[13] L. T. and L. X. Optimal mean reversion trading: Mathematical analysis and practical applications. *World Scientific Publishing Co*, 2016.

[14] E. R. Weintraub. Neoclassical economics. *The Concise Encyclopedia Of Economics*, 2007.