Online intra-day portfolio optimization using regime based models

A Masters Thesis by Sara Fremlin

February 2019

Abstract

In this thesis model predictive control (MPC) is used to dynamically optimize a portfolio where the data is sampled every 5 minutes. Previous research has shown how MPC optimization applied to daily sampled financial data can generate a portfolio that exceeds the value of standard portfolio strategies such as Strategic asset allocation. MPC has been found to have a computational advantage when return forecasts are updated every time a new observation becomes available. A two-state Hidden Markov Model with time varying parameters is used to forecast the financial return of a market index. The portfolio optimization is performed using both single period and multi-period forecasts where the only other asset is a zero interest rate cash account. Transaction costs are included to better reflect market conditions and to address estimation errors in the forecasts. The MPC portfolios are found to outperform a buy and hold strategy in the market index, displaying both higher returns and lower risk. The multi-period portfolios display lower returns and similar risk to the single period portfolio while having a smaller turnover. This led to the conclusion that the two-state Gaussian HMM provides sub par multi-period forecasts on the 5 minute sampled market index. The forecasting method is found to be very sensitive to the manual choice of hyperparameters.

Acknowledgements

I would like to thank Johanna Lägnert for a crash course on the Bloomberg Terminal. My fellow master thesis writers in Gupta, for countless interesting discussions and exchange of ideas over fika. My family and friends for all the generous support received during the entire process.

Last, but not least, I would like to express my sincere gratitude to my supervisor, Professor Erik Lindström. Without his guidance and input, this thesis would not have been been possible.

Nomenclature

- SAA - Strategic asset allocation
- DAA - Dynamic asset allocation
- HMM - Hidden Markov Model
- MPC - Model predictive control
- SPO - Single period optimization
- MPO - Multi-period optimization
- LL - log-likelihood
- iid - Independent and identically distributed
- B&H - Buy and Hold
- BP - Basis point

Contents

1 Introduction

1.1 Background

In the world of finance, investors aim to maximize returns on investments, often while also considering the investments risk profile. Almost every investor has a different opinion on what this relationship of risk versus return should be. Modern investors are faced with a myriad of financial assets that they can buy and sell, all with different risk profiles. This makes the choice of selecting the optimal combination of assets very difficult. The collection of assets an investor holds is called a portfolio. Many different frameworks and strategies have been developed to optimally solve this portfolio choice problem. Markowitz (1952) introduced such a framework called Modern Portfolio Theory, or mean-variance analysis, where the expected returns for a collection of assets are maximized for any given risk.

Investor are not only faced with the choice of what to buy, but also when and how often they need to reevaluate these choices. Many investors apply a single period strategic asset allocation (SAA) strategy, where portfolio weights are determined by investor's risk appetite according to variations of the mean-variance framework, and then periodically re-balanced. The period in between re-balancing is based on a balance between the investment horizon and adaptivity to market conditions. Pension funds, for example, have a long investment horizon of 10 years or longer, however the portfolios are often re-balanced monthly or quarterly. The challenge with SAA is that the market behaviour can shift between periods, and these new behaviours tend to persist for several periods after a shift (Ang & Timmermann 2012). These different periods of market behavior are called regimes and the shifts between regimes correspond to changes in economic policy, regulation or other economic factors that affect asset prices. The flawed aspect of the SAA is apparent in the case of a large downturn of the market between the re-balancing periods. Even when the value of the portfolio drops, no action is taken to improve the situation until the next re-balancing period. When the re-balancing period arrives, the portfolio is adjusted back to the determined static weights. It is apparent that in a world where the future of the market is uncertain, SAA might not be the optimal strategy.

An alternative to the static strategy is to instead use a dynamic strategy. The dynamic strategy enables portfolio weights to continuously be adjusted as new information arises from the market. This creates the opportunity to take advantage of beneficial market regimes and reduce the impact of adverse regimes. When the market shifts between different regimes, a dynamic asset allocation (DAA) strategy is shown to be a more beneficial strategy than the traditional SAA (Sheikh & Sun 2012). The strength of the dynamic strategy relies on using an appropriate model to predict the new information of the market. Regime switching models, such as the Hidden Markov Model (HMM), have been shown to be able to capture the tendency of the financial market to abruptly change its behaviour. They are also shown to capture many of the stylized facts of financial time series, such as heteroscedasticity, skewness and leptokurticity (Ang & Timmermann 2012). In many studies, DAA strategies based on regime switching models have been shown to be a successful approach to portfolio management (Bulla et al. 2011, Nystrup et al. 2018, 2017).

When investors trade, they need to consider much more than just returns and risks when making decisions. Trading is not a free venture, investors incur transaction costs when they buy and sell assets. Even the large institutional trades have costs associated with trading. It is intuitive to reflect this when optimizing portfolio weights. This can be done by including certain constraints, such as holding and trading costs. Not only does this better reflect the real world of trading, adding costs to trades can reduce the errors of forecasted returns by punishing sub-optimal trades. Model predictive control (MPC) was proposed as an approach to solve the stochastic portfolio optimization problem that includes these kind of constraints (Boyd et al. 2017, Herzog et al. 2007). Herzog et al. (2007) concluded that MPC is a suboptimal control strategy for stochastic systems that uses new information advantageously as it is more computationally efficient than the more commonly used stochastic programming models. This makes it an advantageous method when combined with an online forecasting method.

Boyd et al. (2017) developed an open source Python package that can perform portfolio optimization with MPC both over a single period (single period optimization, SPO) and over multiple periods (multi-period optimization, MPO). For the multi-period case, the convex optimization is performed over several periods simultaneously, but only the trade of the current period is performed. This means when trading costs are included, the result of the multi-period case will differ from the single period. This is because the MPO does not only take into consideration the trades in this period, but also how that trade will affect any futures trades over a time horizon. For example, while the SPO might suggest that it is optimal to go long an asset, the MPO looks ahead in time to see if it is still optimal or if the position has to be unwound while taking into account the costs to unwind the position.

The MPC is reliant on the choice of an appropriate model to forecast the returns and risks of the different assets in the portfolio. The Hidden Markov Model (HMM) is an example of a regime switching model that is able to replicate many of the stylistic facts of financial time series (Ang & Timmermann 2012, Rydén et al. 1998, Nystrup et al. 2015). In the HMM, the probability distribution of an observation depends only on the states of an unobserved Markov Chain. The states in the underlying chain can be viewed as representing the different economic regimes that display different market behaviours. Because asset price dynamics are time varying, it is intuitive to choose an adaptive model. The adaptive model will have time varying parameters that allow the properties of the model to change over time and adapt to new market behavior. The forecasts using an HMM will be mean-reverting, changing mainly based on when the regime probabilities change. This means that anticipating the direction of regime change matters more than perfectly estimating the parameters of the model (Sheikh & Sun 2012). The portfolio allocation in the MPC will thus be decided depending on the probability of a certain underlying regime in the HMM and re-balanced when the probability of reaching a different regime in the model is high enough. Therefore, it is of great importance to choose an appropriate parameter estimation method that can correctly estimate the regime probabilities in the HMM.

1.2 Aim

Empirical studies on regime based dynamic asset allocation have mainly been focused on daily or monthly asset returns. Nystrup et al. (2018, 2017) successfully implemented applying HMM forecasts to MPC portfolio optimization using daily asset returns. The sampling interval of the data often reflects the investors investment horizon. DAA however, provides the opportunity to react to changes in the market and rebalance the portfolio if needed. This opportunity may be taken advantage of by using a shorter data sampling interval.

The aim of this thesis is to explore the potential application of MPC to optimize a portfolio using intra-day asset prices. The asset price dynamics are modeled with a two state HMM with Gaussian components.

The results are evaluated by comparing the performance of different MPC implementations to a standard buy and hold strategy.

2 Theory

2.1 Hidden Markov Model

The model used to forecast asset returns and infer the hidden regimes of the market is a Hidden Markov Model (HMM) with Gaussian components. We let $\{O_t\}_{t=1}^T$ be a stochastic process and $\{o_t, t = 1, .., T\}$ its realization. In a HMM, the probability distribution of O_t depends only on the states of a hidden first order Markov chain S_t. While the realizations o_t and the process O_t are directly observable the underlying chain S_t is not.

A sequence of discrete random variables $\{S_t\}_{t=0}^\infty$ is said to be a first order Markov chain if for all $t \geq 0$ it satisfies

$$\Pr(S_{t+1} \mid S_t, ..., S_1) = \Pr(S_{t+1} \mid S_t) \tag{2.1}$$

The conditional probabilities

$$\Pr(S_{t+1} = j \mid S_t = i) = \gamma_{ij} \tag{2.2}$$

are called transition probabilities. The transition probability matrix $\mathbf{\Gamma}$ is the matrix with elements $(i, j) = \gamma_{ij}$. The distribution of the initial states is denoted $P(S_1 = i) = \delta_i$. The chain has a stationary distribution $\boldsymbol{\pi}$, if $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{\Gamma}$ and $\boldsymbol{\pi}\mathbb{1} = 1$. The process $\{Y_t\}_{t=1}^T$ is called an m-state HMM if the chain S_t has m states. A two state Markov chain will have a transition probability matrix of the form

$$\mathbf{\Gamma} = \begin{bmatrix} \gamma_{11} & 1-\gamma_{11} \\ 1-\gamma_{22} & \gamma_{22} \end{bmatrix} \tag{2.3}$$

The HMM can be summarized as

$$\begin{cases} \Pr(O_t \mid O_{t-1}, ..., O_1, S_t, ..., S_1) \\ \Pr(S_t \mid S_{t-1}, ..., S_1) \end{cases} = \begin{cases} \Pr(O_t \mid S_t) \\ \Pr(S_t \mid S_{t-1}) \end{cases} \tag{2.4}$$

It is worth emphasizing that the hidden states form a first order Markov chain. Given that the current state is known, the observed process is independent of the previous historical observations. Figure (1) illustrates the concept of the HMM with hidden underlying states and an observable process.

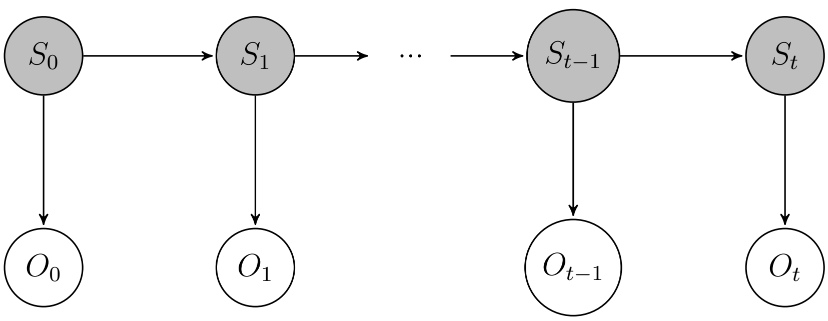

Figure 1: Figure illustrating the concept of a Hidden Markov Model. S represents the hidden underlying states and O represents the observable process

The observed process in the HMM will be the log-returns of the asset, which are assumed to follow a Gaussian conditional distribution

$$o_t|S_t \sim N(\mu_{S_t}, \sigma_{S_t}) \tag{2.5}$$

2.2 Parameter estimation

Because the underlying process of the asset prices are assumed not to be constant, an adaptive parameter estimation method is preferred. A common approach to estimating the parameters in the HMM is the maximum likelihood method. The parameters are estimated online by maximizing the weighted log likelihood (LL) function

$$\hat{\theta}_T = \underset{\theta}{\operatorname{argmax}} \sum_{t=1}^{T} w_t \log \Pr(O_t|O_{t-1}, ..., O_1, \theta) = \underset{\theta}{\operatorname{argmax}}\ \tilde{l}_T(\theta) \tag{2.6}$$

Where $\theta = \{\mu_1, \mu_2, \sigma_1, \sigma_2, \gamma_{11}, \gamma_{22}\}$ are the different parameters to be estimated in the model. When $w_t = 1$ equal weight is put on all previous observations in the calculation of the LL, corresponding to traditional maximum likelihood. The estimator can be made adaptive by changing the weights or the "importance" that older observations will have on the LL. One approach to weighting the observations is to use exponential weights $w_t = f^{T-t}$. $f \in (0, 1)$ is called the forgetting factor, and is determined by the effective memory length

$$N_{\mathit{eff}} = \frac{1}{1-f} \tag{2.7}$$

The choice of the efficient memory length will impact how adaptive the estimation method is. Too small and the parameter estimations will be noisy, too large and the estimations will not be adaptive enough.

By maximizing the second order Taylor expansion of $\tilde{l}_t(\theta)$ a recursive estimator for the parameters is obtained.

$$\hat{\theta}_t = \hat{\theta}_{t-1} - \left[\nabla_{\theta\theta}\tilde{l}_t\left(\hat{\theta}_{t-1}\right)\right]^{-1} \nabla_\theta \tilde{l}_t\left(\hat{\theta}_{t-1}\right) \tag{2.8}$$

where $\mathrm{H}_t = \nabla_{\theta\theta}\tilde{l}_t(\hat{\theta}_t)$ is the Hessian and $\nabla_\theta \tilde{l}_t(\hat{\theta}_t)$ is the score function. To progress with the parameter estimation, a good online estimator is needed for the score function and the Hessian.

Lystig & Hughes (2002) introduced an algorithm that exactly computed the score vector and the log likelihood in one pass. The algorithm is derived from the forward backward algorithm. The method to calculate the score function uses an algorithm to compute the log likelihood.

We let $p_i(o_t) = \Pr(O_t = o_t|S_t = i, \theta_t)$ be the probabilities of an observation given that the state is known. In this case, the Gaussian density. The algorithm of Lystig-Hughes is initialized by calculating the joint distribution of the first time point $Pr(O_1, S_1 = j)$

$$\lambda_1(j) = \Pr(O_1, S_1 = j) = p_j(o_1)\delta_j \tag{2.9}$$

where δ_j is the initial distribution of the hidden states as defined above. It then proceeds by calculating the scaled forward probabilities $\lambda_t(j)$

$$\lambda_t(j) = \Pr(O_t, S_t = j|O_{t-1}, .., O_1) = \sum_{i=1}^{s} [\lambda_{t-1}(i)\, p_j(o_t)\, \gamma_{ij}] \times (\Lambda_{t-1})^{-1} \tag{2.10}$$

where $\Lambda_t = \Pr(O_t|O_{t-1}, ..., O_1) = \sum_{j=1}^{s} \lambda_t(j)$. This leads to the log likelihood $l_T(\theta)$ being able to be expressed as

$$l_T = \sum_{t=1}^{T} \log(\Lambda_t) \tag{2.11}$$

The derived expressions for the log likelihood are now used in an algorithm to compute the score function. The algorithm is initialized with the first component of the score recursion

$$\psi_1(j, \theta) = \frac{\partial}{\partial\theta}\Pr(O_1 = o_1, S_1 = j) = \left[\frac{\partial}{\partial\theta}p_j(o_1)\right]\delta_j + p_j(o_1)\left[\frac{\partial}{\partial\theta}\delta_j\right] \tag{2.12}$$

The partial derivatives $\frac{\partial}{\partial\theta}p_j(o_t)$ of the Gaussian density can be explicitly computed as

$$\frac{\partial}{\partial\mu_j}p_j(o_t) = \frac{\partial}{\partial\mu_j}\phi(\mu_j, \sigma_j^2) = \frac{o_t - \mu_j}{\sigma_j^2}\,\phi(\mu_j, \sigma_j^2) \tag{2.13}$$

$$\frac{\partial}{\partial\sigma_j}p_j(o_t) = \frac{\partial}{\partial\sigma}\phi(\mu_j, \sigma_j^2) = \left(\frac{(o_t - \mu_j)^2}{\sigma_j^3} - \frac{1}{\sigma_j}\right)\phi(\mu_j, \sigma_j^2) \tag{2.14}$$

The algorithm proceeds with calculating the scaled derivatives of the joint distribution

$$\begin{aligned}\psi_t(j, \theta) &= \frac{\frac{\partial}{\partial\theta}\Pr(O_t, ..., O_1, S_t = j)}{\Pr(O_{t-1}, ..., O_1)} \\ &= \sum_{i=1}^{s}\Bigg(\psi_{t-1}(i, \theta)p_j(o_t)\gamma_{ij} + \lambda_{t-1}(i)\left[\frac{\partial}{\partial\theta}p_j(o_t)\right]\gamma_{ij} \\ &\quad + \lambda_{t-1}(i)p_j(o_t)\left[\frac{\partial}{\partial\theta}\gamma_{ij}\right]\Bigg) \times (\Lambda_{t-1})^{-1}\end{aligned} \tag{2.15}$$

The scaled derivatives $\psi_t(j, \theta)$ can then be used to express the score function as

$$\begin{aligned}
\frac{\partial}{\partial\theta}\tilde{l}_T &= \sum_{j=1}^{s} \frac{\frac{\partial}{\partial\theta}\Pr(O_T, ..., O_1, S_T = j)}{\Pr(O_{T-1}, ..., O_1)} \times \frac{1}{\Pr(O_T|O_{T-1}, ..., O_1)} \\
&= \frac{\sum_{j=1}^{s} \psi_T(j, \theta)}{\Lambda_T} \\
&= \frac{\Psi_T(\theta)}{\Lambda_T}
\end{aligned} \tag{2.16}$$

For the adaptive estimation method, the expression of the score vector is adapted to the weighted form

$$\frac{\partial}{\partial\theta}\tilde{l}_{T,H} = w_{T-H+1}\frac{\Psi_{T-H+1}(\theta)}{\Lambda_{T-H+1}} + \sum_{t=T-H+1}^{T} w_t \left(\frac{\Psi_t(\theta)}{\Lambda_t} - \frac{\Psi_{t-1}(\theta)}{\Lambda_{t-1}}\right) \tag{2.17}$$

To estimate the inverse Hessian we use the recursive estimation

$$H_t = \lambda H_{t-1} + h_t h_t^T \tag{2.18}$$

where

$$h_t = \nabla_\theta \tilde{l}_t(\theta_{t-1}) \tag{2.19}$$

The recursive estimation for the inverse is acquired by applying the Sherman-Morrison matrix inversion lemma

$$(A + BCD)^{-1} = A^{-1} - A^{-1}B(C^{-1} + DA^{-1}B)^{-1}DA^{-1} \tag{2.20}$$

The inverse Hessian is then directly recursively estimated by

$$H_t^{-1} = \frac{1}{\lambda}\left(H_{t-1}^{-1} - \frac{H_{t-1}^{-1}h_t h_t^T H_{t-1}^{-1}}{\lambda + h_t^T H_{t-1}^{-1} h_t}\right) \tag{2.21}$$

The approximations for the Hessian and the score function are then inserted into the adaptive recursive estimator

$$\hat{\theta}_t \approx \hat{\theta}_{t-1} + a \cdot H^{-1}\nabla_\theta \tilde{l}_t(\hat{\theta}_t) \tag{2.22}$$

where a is a tuning constant selected to optimize the rate of convergence.

The filter probabilities $\alpha_{t|t}(j) = P(S_t = j|O_t, ..., O_1)$ of being in a certain state can be expressed as

$$\alpha_{t|t}(j) = \Pr(S_t = j|O_t, ..., O_1) = \frac{\Pr(O_t, S_t = j|O_{t-1}, .., O_1)}{\Pr(O_t|O_{t-1}, ..., O_1)} = \frac{\lambda_t(i)}{\Lambda_t} \tag{2.23}$$

Finally, the k-step forward probabilities $\alpha_{t+k|t}(j) = P(S_{t+k} = j|O_t, ..., O_1)$ are obtained by applying the Chapman-Kolmogorov equation

$$\boldsymbol{\alpha}_{t+k|t}^T = \boldsymbol{\alpha}_{t|t}^T \boldsymbol{\Gamma}_t^k \tag{2.24}$$

2.3 Transforms

To ensure convergence of the recursive estimator in equation 2.8, transforms are applied on constrained parameters. A logarithmic transform is applied to variances and the logistic transform is applied to the transition probabilities. The logarithmic transform is

$$\begin{aligned} f(x) = \log(x) \in \mathbb{R} \quad x \in [0, \infty] \\ f^{-1}(y) = e^y \quad f^{-1'}(y) = f^{-1}(y) \end{aligned} \tag{2.25}$$

The logistic transform is

$$\begin{aligned} f(x) = \text{logit}(x) = -\log(x^{-1} - 1) \in \mathbb{R} \quad x \in [0, 1] \\ f^{-1}(y) = \text{logistic}(y) = \frac{1}{1 + e^{-y}} \quad f^{-1'}(y) = f^{-1}(y)(1 - f^{-1}(y)) \end{aligned} \tag{2.26}$$

2.4 Forecasting

For the m-state HMM $\{O_t\}_{t=1}^T$ defined in (2.4) the marginal distribution of O_t can be derived when the probability of being in a state is known, given some information and the marginal distribution of the observed process conditioned on that state (Frühwirth-Schnatter 2006)

$$p(o_t|\mathcal{F}) = \sum_{i=1}^{m} P(S_t = i|\mathcal{F})P(O_t = o_t|S_t = i) \tag{2.27}$$

Given the conditional distribution $p(o_t|S_t = i)$ has the density $p(o_t|\boldsymbol{\theta}_i)$, the unconditional distribution of O_t is obtained as a mixture distribution

$$p(o_t|\mathcal{F}) = \sum_{i=1}^{m} p(o_t|\boldsymbol{\theta}_i)\alpha_i \tag{2.28}$$

The first and second moments of the mixture distribution are

$$\mu = \sum_{i=1}^{m} \mu_i \alpha_i \tag{2.29}$$

$$\sigma^2 = \sum_{i=1}^{m} (\mu_i^2 + \sigma_i^2)\alpha_i - \mu^2 \tag{2.30}$$

where μ_i, σ_i are respectively the forecasted first and second moments conditioned on state $S_t = i$. α_i is the forward state probabilities i.e. the probability of being in state $S_t = i$. The estimated parameters are assumed to stay constant in the multi-period forecasts since no model is implemented to explain their progression. The first and second moment of the mixture distribution are calculated based on the log returns of the asset price. The log returns are calculated by

$$o_t = \log(P_t) - \log(P_{t-1}) \tag{2.31}$$

where P_t is the asset price. For the purpose of optimizing the portfolio we need the forecasted mean and variance of the returns. The returns in each state are assumed to be iid with a log-normal distribution.

$$\log(1 + r_t) \sim N(\mu, \sigma^2) \tag{2.32}$$

where μ an σ^2 are respectively the mean and variance of the log-returns. This gives the expectation and variance for the simple returns

$$\mathbb{E}[r_t] = \exp\left(\mu + \frac{\sigma^2}{2}\right) - 1 \tag{2.33}$$

$$\mathbb{V}[r_t] = \left(\exp\left(\sigma^2\right) - 1\right)\exp\left(2\mu + \sigma^2\right) \tag{2.34}$$

Because the forecasts are based on a stationary Markov Chain, they will be mean reverting. As the k-step forward probabilities approach the stationary distribution of the chain, the forecasts will approach a stationary value as well. This means that forecasting beyond the convergence rate of the chain will be futile as the forecasts will not change.

2.5 Portfolio selection

2.5.1 Model predictive control

In multi-period portfolio optimization the goal is to optimize the total portfolio value over a certain planning horizon K. Given that the future total value of the portfolio is unknown, the optimization problem is formulated as a stochastic control problem. The formulation of the multi-period portfolio optimization as a stochastic control problem follows from Nystrup et al. (2018) version inspired by Boyd et al. (2017). The goal is to maximize the expectation of the total value of the portfolio V_T over the horizon K, subject to cost penalties $\psi(h_t, u_t)$, based on portfolio holdings $h_t \in \mathbb{R}^n$ and value of trades $u_t \in \mathbb{R}^n$

$$\underset{h}{\text{maximize}} \quad \mathbb{E}\left[V_{t+K} - \sum_{t=\tau}^{t+K-1} \psi_\tau(h_\tau, u_\tau) \,\middle|\, \mathcal{F}_t \right] \tag{2.35}$$

The post trade portfolio is then defined as

$$h_t^+ = h_t + u_t \quad t = 0, ..., T-1 \tag{2.36}$$

where $(h_t)_i < 0$ implies a short position on asset i and $(u_t)_i > 0$ means an asset bought. Trades are assumed to be executed at the end of each holding period. Provided that the penalty function is convex, this formulation of the stochastic optimization problem creates a convex objective function, ensuring that a unique solution exists. The total portfolio value V_T is assumed to be a stochastic variable subject to the returns of the assets in the portfolio. Constraints can be included on the holdings to reflect investors preferences

$$h_t^{min} \leq h_t \leq h_t^{max} \tag{2.37}$$

where the constraint $0 \leq h_t$ represents a long-only portfolio. It is natural to assume that an investor does not have access to infinite amount of cash to go long on assets, thus an intuitive approach is to add a self-financing condition to the portfolio

$$\mathbf{1}^T u_t + \kappa^T |u_t| \leq 0 \qquad t = 0, .., T-1 \tag{2.38}$$

That is, that the total proceeds from purchases and sales has to be less or equal to the total transactions cost of performing the trades. Here, κ is a vector of transaction costs that can be taken directly from market information or estimated. For optimization purposes the equality is replaced by a convex relaxation. This can be seen as allowing funds to be removed from the portfolio.

Using MPC the stochastic optimization problem is reformulated as a deterministic problem by expressing the unknown expected total portfolio value in terms of weighted forecasted returns $\hat{r}_{\tau|t}, \tau = t+1, ..., T+H$.

$$\begin{array}{ll} \text{maximize} & \sum_{\tau=t}^{t+K-1} \left(\hat{r}_{\tau|t}^T w_\tau - \psi_\tau(w_\tau) \right) \\ \text{subject to} & \mathbf{1}^T w_\tau = 1, \quad \tau = t+1, ..., t+K \end{array} \tag{2.39}$$

where $\psi_\tau(w_\tau)$ is the penalty function for the costs incurred trough trading and holding.

2.5.2 Trading aversions

The general penalty cost function from (2.39) can be partitioned into different functions to represent certain common trading aversions such as holding costs, transaction costs and risk aversion.

There are many different risk measures used in portfolio optimization in literature. Any convex risk measure could be implemented in the framework of this thesis and many are explicitly implemented in Boyd et al. (2017).

Given that Σ_t is the estimated covariance matrix of the forecasted stochastic returns r_t, the traditional quadratic risk measure of the portfolio is

$$\psi_\tau^{risk}(w_t) = \rho \cdot w_t^T \Sigma_t w_t \tag{2.40}$$

where ρ represents a risk aversion parameter to be tuned in the MPC. In the case of the two asset world where one asset's return is the risk free rate, the covariance matrix will have zeroes in the last row and column since the risk free rate is known. This risk measure combined with the objective function (2.39) corresponds to adjusting the portfolio weights according to the classical mean-variance criterion by Markowitz (1952).

Portfolio optimization based on the mean-variance criterion is known to be sensitive to the estimation errors that occur when forecasting expectations and variances. To compensate for these errors and create trading aversion in the algorithm, a transaction cost is included.

$$\psi_\tau^{trade}(w_t) = \kappa \cdot |w_t - w_{t-1}| \tag{2.41}$$

The transaction cost in this thesis is a scalar. It could be easily replaced by a vector with different values for each asset. This would represent the real world of trading better. For example, trading cash incurs minimal to no transaction costs and trading more illiquid assets might have higher costs than commonly traded assets. Higher transaction costs for certain assets could also represent an aversion towards trading that specific asset in the MPC.

2.5.3 Performance metrics

Certain metrics are used to evaluate the performance of the portfolio. The portfolio return is defined as the average return over the period

$$R_p = \frac{1}{T} \sum_{t=0}^{T} R_t \tag{2.42}$$

The portfolio risk is expressed with the standard deviation

$$\sigma_p = \sqrt{\frac{1}{T} \sum_{t=0}^{T} (R_t - R_p)^2} \tag{2.43}$$

The Sharpe ratio, or information ratio, of the portfolio is a method to evaluate the variance-normalized returns

$$SR = \frac{R_p}{\sigma_p} \tag{2.44}$$

The maximum drawdown is a measure of the largest drop from a peak to a valley and an indication of downside risk. We denote the maximum up until time T as

$$M_T = \max_{t \in (0,T)} P_t \tag{2.45}$$

the drawdown D_t is now defined as the difference of price compared to the maximum

$$D_t = \frac{M_t - V_t}{M_t} \tag{2.46}$$

Maximum drawdown is then defined as

$$MDD_t = \max_{t \in (0,T)} D_t \tag{2.47}$$

The portfolio turnover is a way to measure the amount and size of trades relative to the total size of the portfolio

$$TO = \frac{1}{T} \sum_t \frac{|u_t|}{V_t} \tag{2.48}$$

3 Method

3.1 Data

The data selected for the asset part of the portfolio was the S&P500 (SPX). The S&P500 was chosen because it reflects the general international market movements and is very liquid. An illiquid asset might cause forecasting issues using the HMM due to missing samples or time intervals where the asset price doesn't change or is not available to trade. An illiquid asset could very well be included in an MPC portfolio, for example by either changing the forecasting model or adjusting transactions costs to penalize trading. The HMM, however, requires the presence of regimes to be able to provide good forecasts. With an illiquid asset the price movements become very small, making regime detection in the HMM very difficult. The historical price data was retrieved from Bloomberg at five minute sampling intervals from 2017-12-12 to 2018-11-26. This consists of roughly one year of five-minute data with n=19453 samples. The sample size was considered large enough for the adaptive parameter estimation methods to initiate correctly and to give an indication if the implementation was satisfactory. As seen in Figure 2, the data illustrates typical behaviour of a financial process that might be appropriate to model with a HMM. There are general trends with periods of increases and decreases. It also displays periods of increased- index price volatility.

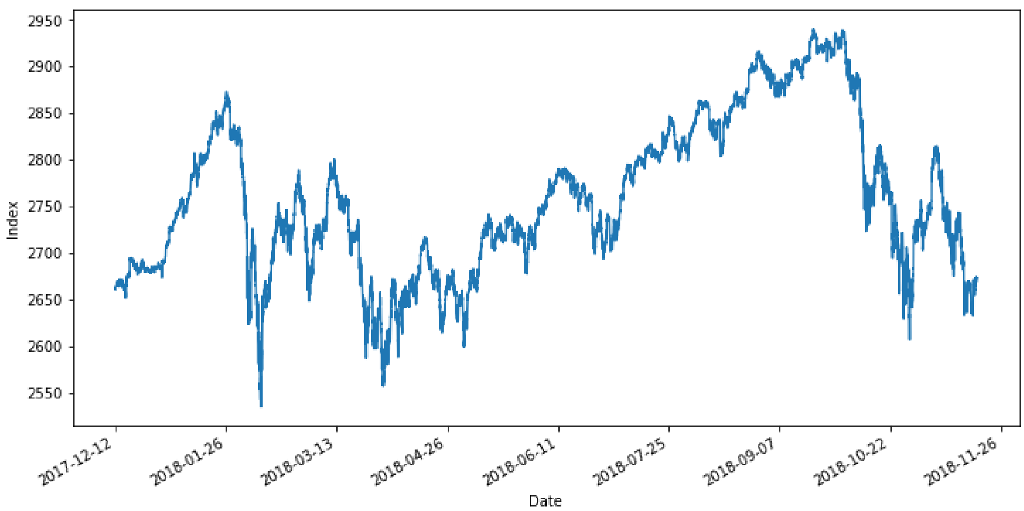

Figure 2: Price process for the S&P500 from 2017-12-12 to 2018-11-26

The data was split in two parts: an in-sample and an out-of-sample period. The in-sample period was used for tuning the HMM hyperparameters, ensuring a good balance between convergence and adaptivity. The in-sample period was the first 13000 samples from 2017-12-12 to 2018-08-02. When the MPC is performed, the full data set was used including the out-of-sample data. Using out-of-sample data in the final tests is beneficial to ensure stability and performance of both the forecasting method and the MPC.

3.2 Implementation

The implementation of the portfolio optimization problem was solved using Python. The MPC was performed by using the CVXPortfolio package for Python (Boyd et al. 2017). The advantage of using the CVXPortfolio package is that forecasting is performed completely separate from the portfolio optimization. This allows any forecasting method of choice to be used in the MPC. If the methods were to be implemented in a real life trading decision, it is straightforward to combine the forecasting and MPC step in one program. This is due to the online parameter estimation method chosen for the forecasts.

Algorithm 1 MPC algorithm

1: Update model parameters based on the most recent observation
2: Forecast asset returns and covariances over a horizon
3: Compute the optimal sequence of trades over that horizon
4: Execute the first trade in the sequence and return to step 1.

The HMM parameter estimation method was initially applied on a simulated two-state HMM to ensure that the code and theory was correctly implemented. The Hidden Markov Model was manually tuned in-sample by choosing appropriate hyperparameters. The hyperparameters for the model were: a, for tuning the adaptivity of the parameter estimates, λ, for tuning the adaptivity of the inverse Hessian estimation, and N_{eff} for tuning the forgetting factor of the adaptive estimation method. The parameter estimation method was tested on simulated data with a non-adaptive method to ensure the parameters converged to reasonable values and the implementation was correct. The adaptive parameter estimation was initialized with manually tested parameters.

Different combinations of MPC hyperparameters were selected to reflect different investment approaches and market situations. The transaction costs were set at 1 basis point (0.0001) and 10 basis points (0.001). 1 basis point is believed to be a realistic transaction cost for high frequency institutional traders. Furthermore, utilizing higher transactions costs can be viewed as a safeguard against forecasting errors. The higher basis point costs will create a kind of barrier that the forecasts need to cross before a trade is considered. The different hyperparameter combinations that were tested in the MPC are represented in Table 1.

Table 1: Different hyperparameter choices for the MPC. κ is the transaction cost. ρ is a risk-averseness parameter. K is the look ahead horizon in the multi-period case and where K=1 is the case of single period optimization.

Hyperparameters	Value
κ	0.0001 , 0.001
ρ	0, 1, 5
K	1, 10, 20, 50, 100

Certain constraints are applied to the portfolio optimization. The portfolio is assumed to be long-only and no leverage is allowed. This means that the weights are not allowed to be negative or exceed 1 ($0 \leq w \leq 1$). This means that shorting is not allowed, nor is borrowing money to buy an asset. The second asset in the portfolio will be risk-free cash with zero interest rate. This means that when equity is sold the cash asset will not increase in value. In a real world situation, the cash asset would often be traded for a risk-free asset with an interest rate, such as Treasuries. For this proof of concept, zero interest is sufficient, as adding interest would simply add higher return when no asset is held. The zero return could easily be exchanged for a risk-free rate if desired. The MPC portfolios are compared to a buy & hold strategy that is 100% in equity. The buy & hold strategy is simply buying a position in the equity and holding until the trading period is over. Holding a certain fixed amount in the risk-free asset, which would be replaced by bonds in most traded portfolios, is a much more common trading strategy than 100% equity. Often portfolio compositions such as a 60%/40% bonds and equity is used. This is commonly done to reduce the risk of the portfolio, while still keeping the potential of higher returns by holding some equity. The comparison of the results of the MPC portfolios to a portfolio that is 100% equity is to evaluate if the portfolios can outperform the market.

Due to the long computation time for backtesting 19000 datasamples in a multiperiod setting, the different lengths of the horizon, K, for the MPO were tested only on the first 5000 samples.

4 Results

The results from implementating and forecasting the HMM using the adaptive parameter estimation methods are presented first, followed by the results of the different MPC versions.

4.1 Hidden Markov Model

The HMM hyperparameters were tuned to a=0.004 and $\lambda = 0.99$. The efficient memory length was selected to be $N_{eff} = 250$. The choice of the efficient memory length is a balance between the parameter estimates adaptivity and stability. If the efficient memory length is too long, the estimates will converge to the traditional maximum likelihood estimation, if it is too short the estimates will have a high variance. Certain hyperparameter choices cause the parameters estimates to not converge and quickly deteriorate to unreasonable estimates. Table 2 displays the results of the HMM hyperparameter selection.

Table 2: Results of the HMM hyperparameter selection

Hyperparameters	Value
a	0.004
λ	0.99
N_{eff}	250

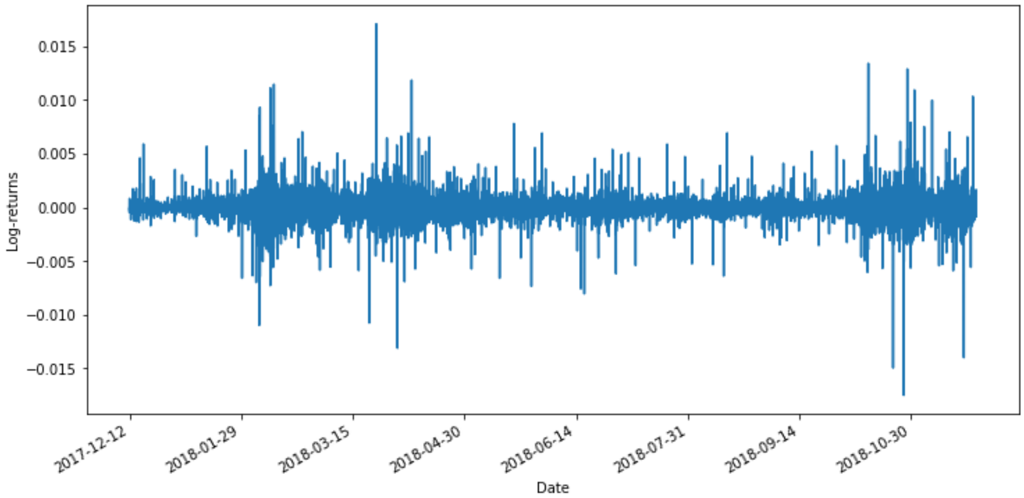

(a) Log returns for the S&P500 from 2017-12-12 to 2018-11-26

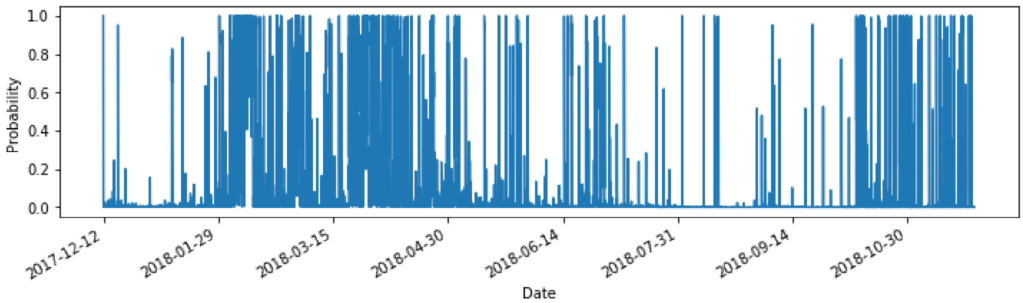

(b) The filter probability of being in state 2 as calculated with the HMM $P(S_t = 2)$

Figure 3: The log-returns and filter probability of being in state 2 (the low return, high variance state as estimated with the HMM)

Figure 4 displays the result of the adaptive parameter estimations on the five minute log returns of the S&P500. The results show that the parameters start fairly stable for the first 3 months and then vary during the rest of the year. Compared to the estimated forward probabilities in Figure 3, the parameters start varying more when the probability of being in the second state starts increasing. The two states consists of one with positive drift and low variance, and one with negative drift and high variance. The probability of remaining in the same state is rarely lower than $\gamma_{11} = 0.995$ for the positive drift state, and rarely lower than $\gamma_{22} = 0.95$ for the second, negative drift state. Both the drift and diffusion of the second state show large jumps and movements during the second half of the year.

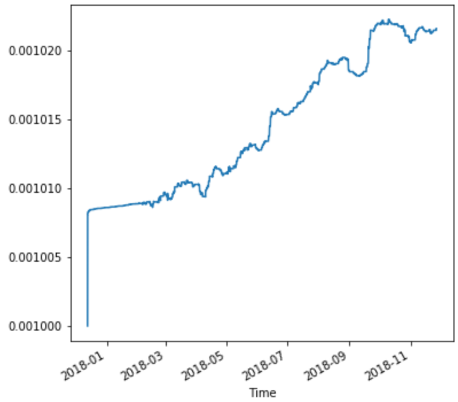

(a) μ_1

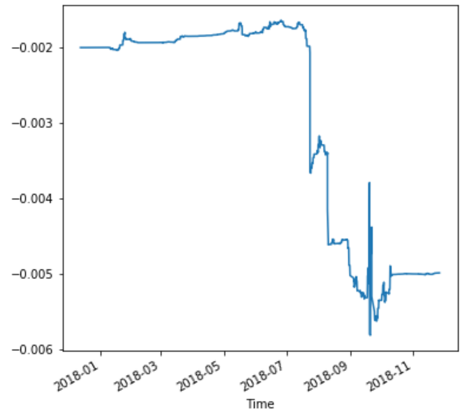

(b) μ_2

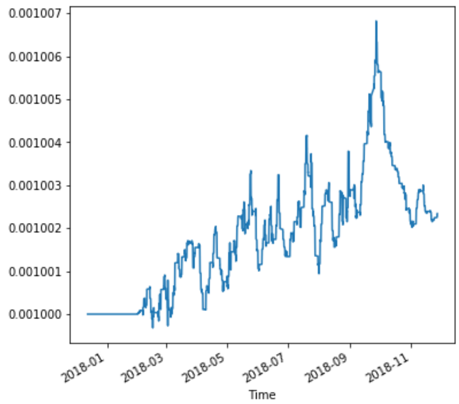

(c) σ_1

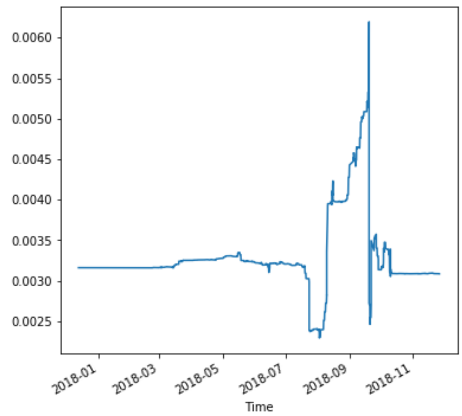

(d) σ_2

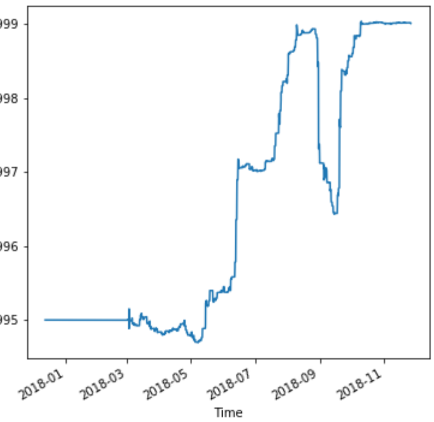

(e) γ_{11}

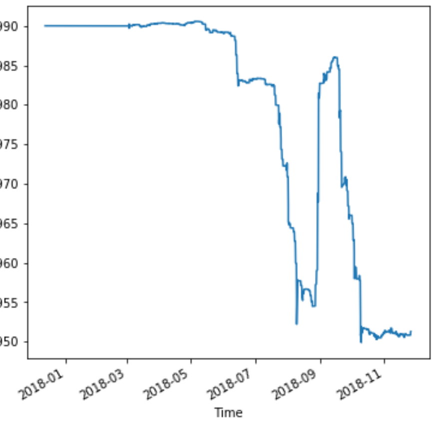

(f) γ_{22}

Figure 4: Adaptive parameter estimates for the two state HMM

4.2 Dynamic portfolio optimization

The results of the single period optimization is described in Table 3. The first case ($\rho = 0$) represents an investor who does not take risk into consideration, only returns. The cases of the more risk-averse investor ($\rho = 1$, $\rho = 5$), i.e. a larger penalty for risk in the optimization, show an increasing turnover and slightly less excess risk the higher the risk averseness parameter. The cases of the low and no risk averseness ($\rho = 0$, $\rho = 1$) show very similar results in terms of portfolio return and excess risk as seen the Sharpe ratios of 0.802 and 0.804 respectively. When a higher transaction cost is introduced ($\kappa = 0.001$) the turnover and return both drastically decrease. Figure (6) illustrates how the high transaction costs affect the trading algorithm.

Table 3: Single period portfolio optimization results for different MPC hyperparameter combinations compared to the Buy and Hold (B&H) strategy. The transaction costs are presented in basis points (BP), where 1BP corresponds to $\kappa = 0.0001$

	ρ=0, 1BP	ρ=1, 1BP	ρ=5, 1BP	ρ=0, 10BP	B&H
Portfolio return (%)	0.837	0.838	0.831	0.028	0.006
Excess risk (%)	1.044	1.043	1.040	0.732	1.586
Sharpe ratio	0.802	0.804	0.799	0.038	0.003
Max. drawdown	2.726	2.726	2.726	6.503	11.879
Turnover (%)	720.641	723.686	721.762	207.578	0

Table 4 displays the results of the MPC using multi-period optimization with different forecast horizons. Turnover is reduced for all cases of the MPO (K > 1) compared to the SPO. The results of the 20 period and 50 period horizon forecasts are very similar and display an increased Sharpe ratio compared to the 10 period horizon. All MPO portfolios outperform the buy & hold Strategy in all metrics. The SPO performs better than all other tested MPO portfolios, displaying higher returns, higher Sharpe ratio and lower drawdown. The risk of the different portfolios is similar for both the SPO and the MPO.

Table 4: Results of the MPO and SPO on the first 5000 samples ($\rho = 0$, $\kappa = 0.0001$).

	K=1	K=10	K=20	K=50	K= 100	B&H
Portfolio return (%)	1.071	0.699	0.705	0.703	0.701	0.183
Excess risk (%)	0.932	0.939	0.939	0.939	0.939	1.586
Sharpe ratio	1.149	0.744	0.751	0.749	0.747	0.110
Max. drawdown	2.015	3.504	3.409	3.421	3.429	11.879
Turnover (%)	665.610	552.661	561.799	562.371	561.429	0

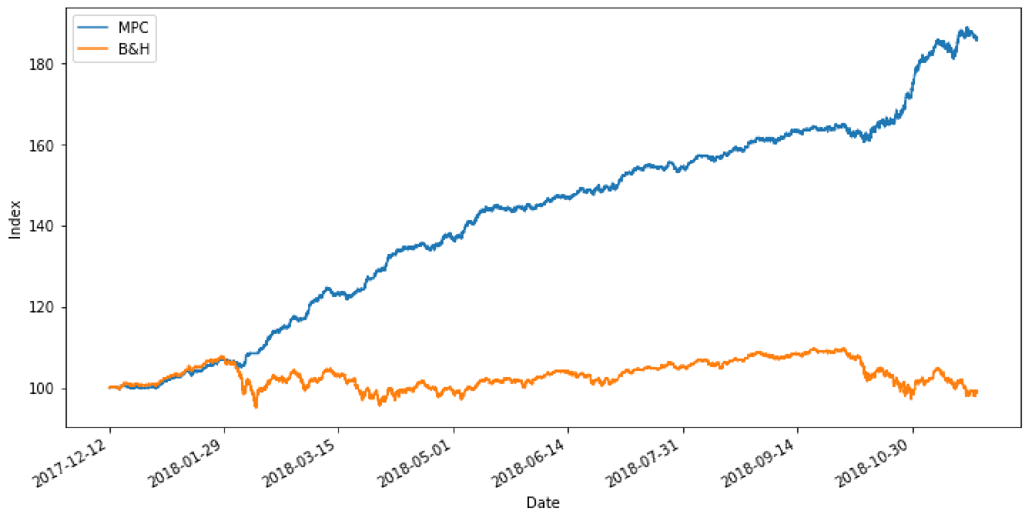

(a) Portfolio optimization results, the single period MPC is compared to a Buy and Hold strategy.

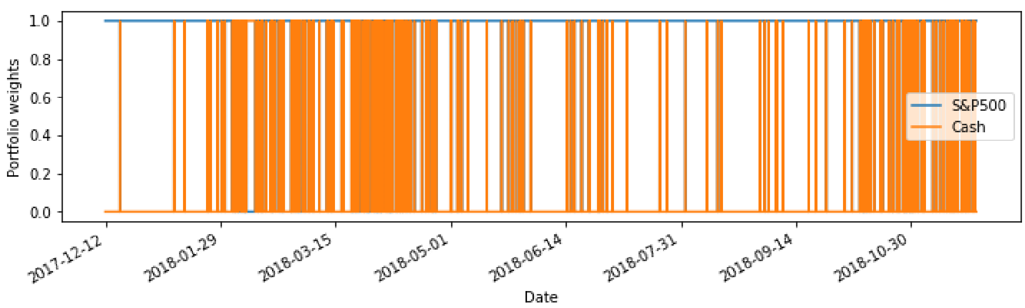

(b) Portfolio weights

Figure 5: Performance of MPC on the full data from from 2017-12-12 to 2018-11-26 (K=1, ρ=0, κ=0.0001).

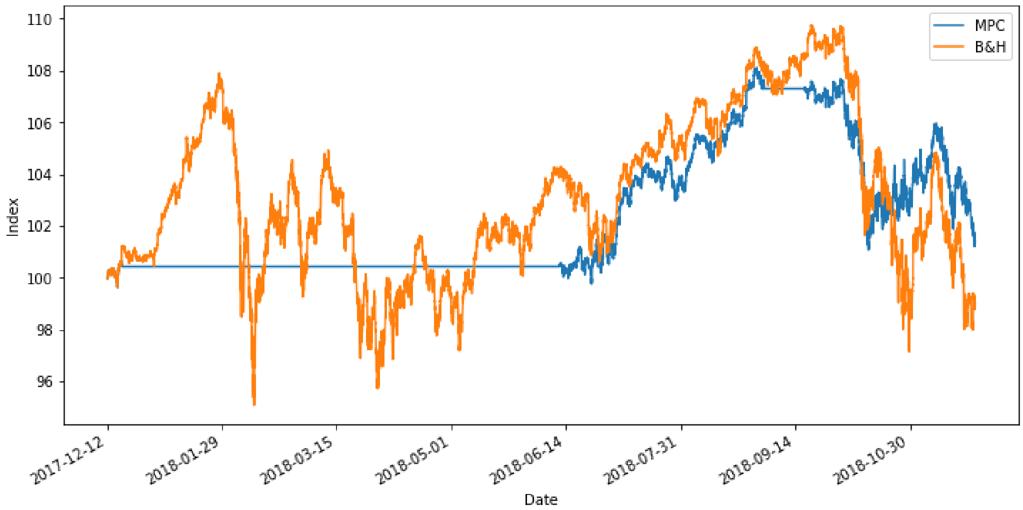

(a) Portfolio optimization results, the MPC is compared to a buy and hold strategy.

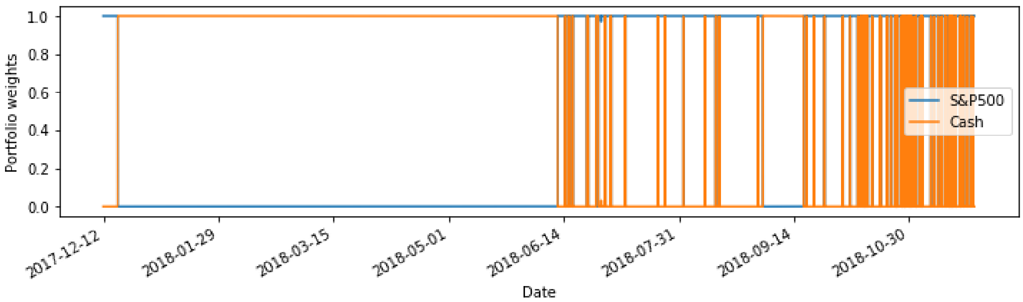

(b) Portfolio weights

Figure 6: Performance of MPC on the full data from from 2017-12-12 to 2018-11-26 (K=1, ρ=0, κ=0.001).

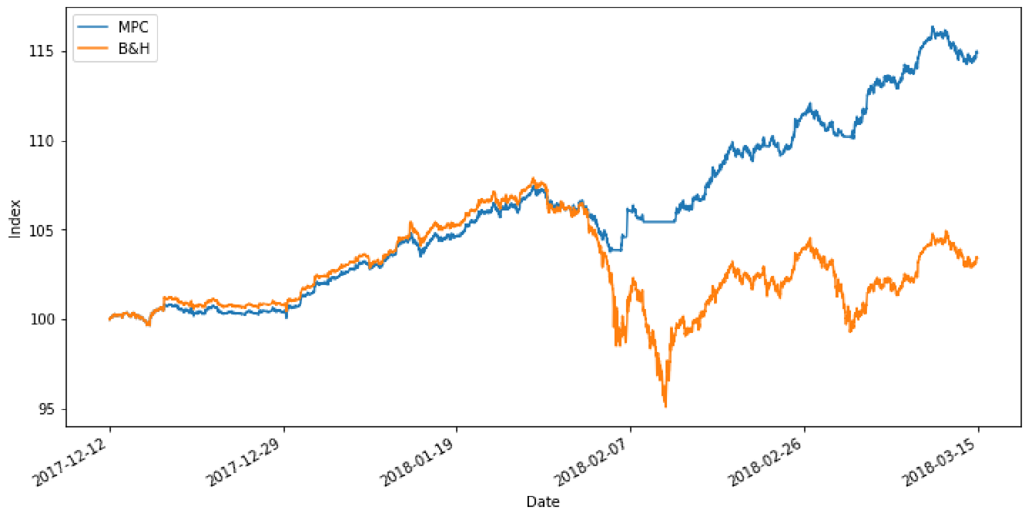

(a) Portfolio optimization results, the MPC is compared to a buy and hold strategy.

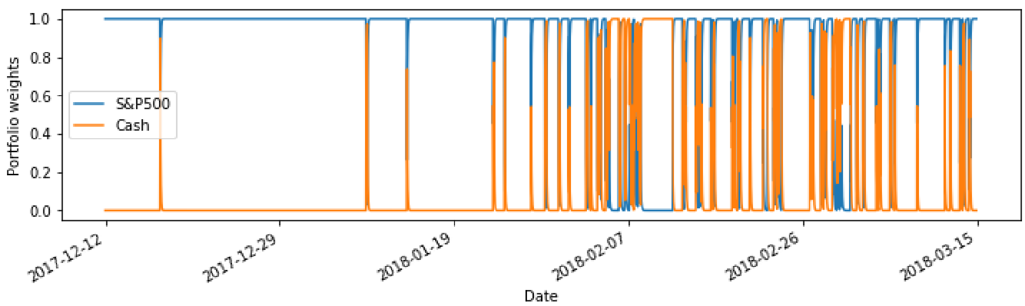

(b) Portfolio weights

Figure 7: Portfolio optimization results for the MPO portfolio (K=20, $\kappa = 0.0001$, $\rho = 0$) over the first 5000 samples.

5 Discussion

The parameter estimates demonstrate a typically expected behaviour of a two state HMM applied to financial time series. The states appear to reflect a period of the equity market when prices generally increase and display lower variance and generally decrease and display higher variance. The probability of staying in the same state does not exceed 0.999, which reflects an expected sojourn time of 1000 periods, or 3.5 days. This indicates that it should not be necessary to test much longer horizons than K=1000 in the MPC. A horizon of K=1000 is very long compared to previous studies and was considered unreasonably long when testing hyperparameters for the MPC. The long sojourn time could indicate that the 5-minute sampling is simply to high frequency for sake of using the forecasts in the MPC. Assuming the HMM properties would stay the same, changing the data sample to every 30-minutes would make the 3.5 day sojourn time be represented in about 167 periods, a much more reasonable horizon for the MPC. The state with a positive mean is more persistent than the state with a negative mean. In absolute terms, the mean of the positive state is also smaller than the mean of the negative state. This describes a longer, generally slow climb of the market which sometimes shifts to rapid declines with high volatility. These results reflect how financial time series and the financial market generally behaves. The model appears to have been able to capture the economic regime shifts and trends previously described in the introduction. In the final quarter of 2018, the market was hit with uncertainty and large downturns as seen in Figure 2. The HMM parameters reflect this behaviour with a large decrease in the negative state's mean and an increase in volatility in both states.

The adaptive HMM parameters occasionally display large jumps. This is especially visible in the volatility estimate for the second state as seen in Figure 4d. In a two month period around September, the estimate shoots up from 0.0025 to 0.006 to later revert. While the parameter estimates converge, this behaviour is not optimal and is a potential issue. This could be caused by the inverted Hessian estimates used in the estimation method. The inverted Hessian occasionally becomes very small and close to indefinite. This leads to the parameter estimation taking a step in an almost random direction. Other Hessian estimation methods could be tested to investigate what is causing the issue. Another method could be to adjust the Hessian when it is close to indefinite, for example by adding a small number on diagonal.

The hyperparameters for the HMM were manually tested. While it worked in the context of testing whether the model was even applicable, it is not intuitively an optimal approach. The different hyperparameters are dependent, this means that a small adjustment for one might mean that another needs to be calibrated as well. This causes a very iterative manual approach whereby one parameter is slightly changed, the estimation algorithm is run, and the whole process is repeated for each parameter until the estimates converge. Manual testing is thus very time consuming and the process ends when the parameters are arbitrarily "good enough". This means that the values of the hyperparameters might not be the optimal combinations for the best forecasts. A different approach to hyperparameter selection could be to test a large number of hyperparameter combinations and optimize for the best combination. However, as the price process is ever changing, there is no guarantee that the optimized hyperparameters would be optimal out of sample. With this in mind, manual selection is not an unreasonable approach.

The choice of five-minute intervals presented some difficulties in the implementation of a forecasting model. Several different forecasting models were investigated before settling on the HMM. The issue was largely the relatively small returns between the short intervals. Even when the HMM was selected, the parameter estimation was not straightforward. Methods such as the adaptive expectation-maximization algorithm was not able to find distinctive regimes. The convergence of the parameter estimations were, in the end, very sensitive to the choice of hyperparameters. Inferior hyperparameters would cause the algorithm to completely deteriorate and stop functioning. This is an indication of the significant sensitivity of the forecasting method. If the method were to be used in a real life trading situation, changes in market behaviour may lead to the algorithm completely failing. Critical failure is often an unacceptable quality of a model. The method could, for example, never be used in its current state for a nuclear power plant. This makes it imperative to monitor the behavior of the hyperparameters in the algorithm and perhaps periodically evaluate and change the values.

The results illustrate how the MPC portfolios consistently outperform the buy and hold strategy for the S&P500 in the sampled time period. The outperformance persists even when the transaction cost is raised up to $\kappa = 0.001$, however it is considerably smaller. The impact of the high transaction costs is displayed in Figure 6. The MPC portfolio reacts by trading less often, and by only holding the asset when the forecasted returns are larger than the high transaction costs. This highlights the close relationship between the forecasts and the decision making in the MPC, and the importance of choosing the forecasting method and penalty functions. The performance of the MPC portfolio is particularly interesting during the last quarter of 2018. The equities market was faced with large downturns and uncertainties causing a period of negative return for the S&P500. As illustrated in Figure 5 the algorithm manages to take advantage of the regime shifts and generates a positive return during this period. This can almost be seen as a stress test for the algorithm, given the period is out of sample, and indicates the advantages of a dynamic trading strategy. While the multi-period results generate a lower return than the SPO strategies, it is not unreasonable to assume that the behaviour MPO would follow the favorable results of the SPO in this uncertain period. Figure 7 illustrates how the MPO MPC manages a positive return during a period of negative returns for the buy and hold strategy.

The different MPO portfolios all display worse results compare to the SPO portfolio. They have worse returns, larger drawdowns and higher Sharpe ratios. Theoretically, the MPO should outperform the SPO if the forecasts are correct, especially if transaction costs are present. The MPO should prevent costly, unnecessary trades and detect long periods of market downturn earlier due to the multi-period forecasts. This is an indication that the HMM model might not be correct, and that the multi-period forecasts are not good enough. Even though the implementation of the HMM was successful in the scope of the thesis, all the MPC portfolios outperformed the buy and hold strategy. An additional argument could be made if it is a good model for the chosen data. The estimated regime probabilities in Figure 3 illustrate a long period of several months before the regime switches that does not coincide with the results of the estimated of the transition probabilities. The transition probability of $\gamma_{11} = 0.995$ would imply an expected sojourn time of 200 periods, much longer than the estimated time spent in the first regime.

The results of the 20, 50 and 100 horizon MPO were very similar and produced more favorable results in terms of a higher Sharpe ratio than the 10-period horizon MPO. The reason that the results are so similar could be due to the persistent states of the underlying Markov Chain of the forecasts. With persistent states, the chain's probabilities converge very slowly towards the stationary distribution, leading to very similar forecasting results over the different horizons.

While all the different versions of MPC hyperparameter combinations outperformed the buy and hold strategy, the optimal combination of parameters may not have been tested. Here, the same argument could be made as with the HMM hyperparameter. Testing manually is time consuming and not necessarily optimal. A combination of hyperparameters could be tested. A coarse search on the risk-return plane could then be performed to find the pareto-optimal combinations that generate the highest return for any given risk level.

The MPC results exhibit a high annual turnover compared to previous studies by Nystrup et al. (2018, 2017). The higher turnover could simply reflect that five minutes is a considerably higher sampling frequency compared to previous research done on daily or monthly sampling periods. For further research it might be interesting to investigate the consequences of using MPC for portfolio optimization on every 10 minutes to hourly data samples. The longer sampling periods might make parameter estimation and forecasting less sensitive to hyperparameters, while maintaining the advantages of intra-day trading periods.

6 Conclusion

The results of this thesis illustrates that MPC portfolio optimization methods can be applied successfully to a five minute sampled financial asset and outperform a buy and hold strategy. The MPC approach results in higher returns, lower risks and lower drawdown than the standard buy and hold strategy, even during times when the asset price is generally decreasing in value. The multi-period MPC reduced the turnover of the portfolio while still providing a higher return and reduced risk compared to the standard strategy. A reduced turnover is favorable when considering the costs of executing the trades. The worse performance of the MPO compared to the SPO indicates that the HMM is not the optimal model for the 5-minute S&P500 samples. The success of the method relies heavily on the chosen forecasting method and hyperparameters. For the purpose of the thesis, a two state HMM with Gaussian components was found to be an acceptable model for the S&P500 over the sampled period, even though the multi-period forecasts did not perform as well as expected. While the implementation of the HMM was successful, it was very sensitive to the manually chosen hyperparameters used in the adaptive parameter estimation algorithm.

Too keep the scope of the thesis reasonable, the focus was on one asset and an interest free cash account. The theory of the thesis is simple to extend to a multi-asset portfolio if covariances between assets can be forecasted. Due to the performance of the MPO, further models to forecast the asset prices should be explored. While the Gaussian components provided a satisfactory model for the purpose of this thesis, log returns are generally leptokurtic, especially on the downside. Thus, introducing a Students-t distributed component might provide better results due to its leptokurtic nature. A three state HMM could also be tested, where the third state could represent a third economic regime of stagnant, low movement price dynamics. If the two-state Gaussian HMM model was to be used, lower frequency samples could be tested, such as hourly or every half hour. A lower sampling frequency might make state detection easier in the HMM due to larger movements in the returns. The efficient memory length might require shorter forecast horizons, shortening the MPO MPC computation time. Since the manual choice of hyperparameters was of large importance, further research could be conducted in simplifying or optimizing the choice of these hyperparameters. This applies both to the HMM and the MPC.

References

Ang, A. & Timmermann, A. (2012), 'Regime changes and financial markets', *Annu. Rev. Financ. Econ.* **4**(1), 313–337.

Boyd, S., Busseti, E., Diamond, S., Khan, R. N., k. Koh, Nystrup, P. & Speth, J. (2017), 'Multi-period trading via convex optimization', *Foundations and Trends in Optimization* **1**(1), 1–76.

Bulla, J., Mergner, S., Bulla, I., Sesboüé, A. & Chesneau, C. (2011), 'Markov-switching asset allocation: Do profitable strategies exist?', *Journal of Asset Management* **12**(5), 310–321.

Frühwirth-Schnatter, S. (2006), *Finite mixture and Markov switching models*, Springer Science & Business Media, pp. 301–318.

Herzog, F., Dondi, G. & Geering, H. (2007), 'Model predictive control for portfolio selection', *Int. J. Theor. Appl. Finance* **10**(2), 203–233.

Lystig, T. C. & Hughes, J. P. (2002), 'Exact computation of the observed information matrix for hidden markov models', *Journal of Computational and Graphical Statistics* **11**(3), 678–689.

Markowitz, H. M. (1952), 'Portfolio selection', *J. Finance* **7**(1), 77–91.

Nystrup, P., Boyd, S., Lindström, E. & Madsen, H. (2017), 'Multi-period portfolio selection with drawdown control', *Annals of Operations Research* pp. 1–27.

Nystrup, P., Madsen, H. & Lindström, E. (2015), 'Stylised facts of financial time series and hidden markov models in continuous time', *Quantitative Finance* **15**(9), 1531–1541.

Nystrup, P., Madsen, H. & Lindström, E. (2018), 'Dynamic portfolio optimization across hidden market regimes', *Quantitative Finance* **18**(1), 83–95.

Rydén, T., Teräsvirta, T. & Åsbrink, S. (1998), 'Stylized facts of daily return series and the hidden markov model', *Journal of applied econometrics* **13**(3), 217–244.

Sheikh, A. Z. & Sun, J. (2012), 'Regime change: Implications of macroeconomic shifts on asset class and portfolio performance', *Journal of Investing* **21**(3), 36.