Erasmus University Rotterdam
Erasmus School of Economics
Master Thesis: Quantitative Finance

Factor Timing Using a Markov-switching Model

Author: Bouwe van der Wal
Student ID number: 458020

Supervisor: dr. H.J.W.G. Kole
Second assessor: dr. R. Lange

Abstract

In this thesis I model the joint process of weekly investment factor returns using a Markov-switching model. Subsequently, I examine the effect of incorporating multiple regimes in a portfolio optimization problem on the investment decisions and performance of a mean-variance investor that manages a portfolio of factors. This is done for a number of investment horizons. I test specifications of a Markov-switching model with different numbers of states, and with explanatory variables predicting either the transition probabilities, or the regime-dependent returns. Variable selection is applied using the LASSO, and a step-wise procedure using the score test. I find that a model with three states, constant regime-dependent returns, and time-varying transition probabilities models the joint process of investment factors most properly. This model identifies regimes based on their volatility. It is able to recognize periods of high volatility out-of-sample, in which it invests more conservatively than a model without regimes. Consequently, it achieves a Sharpe ratio that is more than 40% greater than that of a model without regimes. Even more, a Markov-switching model achieves Sharpe ratios more than twice as large as those from a no-regime benchmark over multi-period investment horizons. It is able to do so by hedging away risks of possible future adverse states. This results in a lower realized volatility compared to a no-regime benchmark.

December 6, 2021

The content of this thesis is the sole responsibility of the author and does not reflect the view of the supervisor, second assessor, Erasmus School of Economics or Erasmus University.

Contents

1 Introduction

Since the first appearance of factors as anomalies that are in conflict with the well known Capital Asset Pricing Model, factors are there to stay in the investment literature. The seminal paper of Fama and French (1993) was the start of wide adaptation of the idea of factors driving equity returns, mainly thanks to the simple definition and construction of the size and value factors, and the empirical evidence provided in the paper. Since then, many more factors have been proposed, but only a few are considered persistent enough to have properly nested themselves in the literature.

While the consensus is that factor premia are persistent over time, their returns are by no means constant. The momentum factor for example has historically exhibited very high returns when a bubble is forming, such as the dot-com bubble. Conversely, the value factor has not shown a large risk premium ever since the financial crisis. From the theory underlying factors, varying performance over shorter horizons however should not come as a surprise: by investing in assets, an investor is exposed to factor risks – that is, the risk that assets pay off little during bad times. How such a bad time is defined depends on the stochastic discount factor, but the most general is a market crash, corresponding to the market factor in the Fama-French three-factor model. In return for their low returns in bad times, in the long-run an investor should be paid a risk premium that generates high returns over the long run (Ang, 2014, Chapter 6).

Thanks to the time-varying nature of factor risk premia, the concept of factor timing becomes attractive for investors. Factor timing is the process of overexposing a portfolio of assets to factors of which it is predicted that they will perform well, while underexposing the portfolio to those which will not. If it is successfully implemented, the investor can capture the risk premia awarded for exposure to bad times, without actually suffering during these bad times.

In practice however, factor timing is hard. This is exactly what is to be expected from the efficient markets theorem; after all, if timing factors was easy, everyone would do it and the premium would vanish.

In this thesis, I will examine what Markov-switching model best describes the dynamics of factor returns, and additionally what these dynamics are. Furthermore, I investigate the effect of incorporating multiple market regimes in a portfolio optimization problem on the decisions and performance of a mean-variance investor that manages a portfolio of factors, possibly giving rise to 'factor timing'. Specifically, I will research the effect of explanatory variables, both in the equations of transition probabilities between regimes, and as predictors of regime-dependent returns: do they improve the model, and do they enhance investment performance?

The dataset used in the research consists of returns on some of the established factors, namely the quality, momentum, value, size and market factor. The returns of the former four are constructed – following Fama and French (1993) – as zero-cost portfolios, resulting in four portfolio return series: small minus big (SMB), high minus low value (HML), quality minus junk (QMJ), and up minus down (UMD). The market factor is the excess market return (MKT). In addition, multiple explanatory variables are considered, observed with a weekly frequency. The data spans a period starting in January 1986 and ending in the last week of December 2020.

A Markov-switching model is used to model the joint dynamics of factor returns. This model identifies multiple regimes, where returns and covariance matrices are dependent on the prevailing regime. To estimate the model, the expectation-maximization (EM) algorithm is used.

There are two reasons why I hypothesize that a Markov-switching model may have good performance in factor timing. The first one is that its structure can capture some of the stylized behaviour of returns, such as: periods with different mean returns discussed earlier; volatility clustering; and time-varying correlations (Ang and Timmermann, 2012). Secondly, earlier works have already established that Markov-switching models can have superior performance in the context of asset allocation over models that do not take multiple regimes into account (Guidolin and Timmermann, 2007; Ang and Bekaert, 2004; Guidolin and Timmermann, 2008). All these papers identify a high-volatility regime. Since mean-variance optimization relies heavily on the covariance matrix, separation of volatility per regime can lead to better assessment of the variance.

Next to a simple Markov-switching model, more complex variations are estimated by letting explanatory variables enter the model, either by predicting regime-dependent returns, or by influencing the transition probabilities between regimes. The main focus is on the latter. I hypothesize that endogenous transition probabilities will lead to better out-of-sample (OOS) prediction of the states – and in turn the returns – compared to a model with constant transition probabilities, because the latest information is incorporated. This hypothesis is driven by papers solely focused on predicting market regimes using explanatory variables. Previous research shows that explanatory variables do a better job in predicting regimes than predicting returns (Chen, 2009), and further that there is evidence of predictability in-sample and OOS (Nyberg, 2013). This is in contrast to earlier findings on return predictability, for which weak evidence is found at best (Welch and Goyal, 2008; Ang and Bekaert, 2007). A model where explanatory variables predict regime-dependent returns is however still estimated, to examine the difficulty of time-varying relations that Bender et al. (2018) identify as a barrier to successful implementation of factor timing. For both types of models with explanatory variables, variable selection is applied: for variables appearing in the transition probabilities, a step-wise procedure based on the score test is used, and for those in the equation for regime-dependent returns, the least absolute shrinkage and selection operator (LASSO) is deployed.

The model fit of all models is evaluated over the full sample using information criteria, and metrics specific to Markov-switching models, such as the regime classification measure (RCM) and a form of a density specification test (DST). To test for individual significance of model estimates, t-tests are conducted.

Next, a buy-and hold mean-variance investor with no-shortselling and borrowing constraints is considered for multiple investment horizons, who uses estimates from the models as inputs for a portfolio optimization problem. To arrive at expected returns and covariance matrices, I use calculations based on the law of iterated expectations for the single-period investment horizon, while a Monte Carlo algorithm is used to obtain multi-period moments. Performance of the portfolios is evaluated using the mean and volatility of portfolio returns, Sharpe ratios, and realized utility. Statistical significance between performance is tested using the Ledoit and Wolf (2008) test of equality of Sharpe ratios. The OOS portfolio performance is tested by training the models on half of the sample, after which an expanding window method is used with 10 succeeding splits in the data; the model is trained on the latest vintage and tested on the following split.

I find that a out of multiple specifications with two, three or four states, a Markov-switching model with time-varying transition probabilities, three states and constant regime-dependent returns models the joint density of factors most closely. The regimes are characterized by either low, medium or high volatility, and transitions between regimes almost always occur through the medium-volatility regime. Volatility in the high-volatility regime is almost 50% higher than that in the low-volatility regime. Also, the size and sign of risk-premia and correlations vary between regimes. The regimes are mostly characterized by their volatility, and to a much lesser extent by their returns. Furthermore, return estimates are not found to be robust, while volatility estimates are.

This model can be used to achieve significantly higher risk adjusted returns than a model without regimes. When applied OOS, it achieves a yearly Sharpe ratio of 1.07 versus 0.739. In addition, its cumulative returns are 40% larger. This superior performance is achieved by its ability to correctly recognize periods of high volatility, and in turn invest more in the riskless asset. It further is better able to predict the prevailing regime than a simple Markov-switching model.

In a multi-period investment setting, superiority of a Markov-switching model over a model without regimes is even more distinct. On all investment horizons, ranging from four weeks to one year, a Markov-switching model achieves a significantly higher Sharpe ratio. Its performance however is best for horizons of four and thirteen weeks, with Sharpe ratios more than twice as large as those from a model without regimes. The superior performance is explained by the model keeping volatility low, which it is able to do by hedging away risks of possible future states.

In the literature on Markov-switching models, this thesis mainly contributes methodologically. I show that the LASSO is easily implemented in existing estimation methods due to similarities in the log-likelihood of a Markov-switching model and a weighted least squares regression, such that each state can be seen as a separate LASSO problem. This however reveals the need for further research on how to efficiently optimize the regularization parameters for each regime. This thesis further reveals that variable selection for the transition probabilities using a step-wise procedure on the score test is not optimal when the sample size is small and some transitions scarcely occur, because statistics may become inflated. This calls for the need to also implement the LASSO for this part of the Markov-switching model. Furthermore, I provide an altercation of the multivariate density specification test of Guidolin and Ono (2006), that can properly apply a transformation to multivariate data to arrive at a univariate statistic. To my knowledge, the original multivariate test was faulty. Also, the limitations of estimation of Markov-switching models with time-varying transition probabilities using the EM algorithm are revealed, showing that estimation becomes impractical when too many variables or states are considered.

This thesis further contributes to the factor timing literature by examining one of the barriers of successful factor timing that Bender et al. (2018) identify, namely varying relations between predictors and factors. It turns out that successful factor timing does not hinge on this phenomenon. Implementing varying returns results in OOS estimates that are equally poor. Instead I show that it is more powerful to use predictors to determine the state or direction of the market. Lastly, this research indicates that explanatory variables can have a very different impact on the expected moments depending on whether they are used to predict regime-dependent returns, or the transition probabilities. Even more, variable selection procedures select a very different set for both purposes. Therefore, future research should not blindly use the same variables for regime prediction as for return prediction.

2 Literature

2.1 Return Predictability

What is generally done in the process of factor timing, is to consider a number of predictor variables that estimate the conditional expected return in future periods and to use the predicted returns to tilt portfolio weights to factors that are expected to perform above average. Predictors that are found to have some explanatory power over factor performance are valuation measures, sentiment (Baker and Wurgler, 2006), macroeconomic variables, and recently factor momentum (Gupta and Kelly, 2019). They show significant factor persistence on a one month basis for a very large set of factors, a particularly interesting finding since this autoregressive nature is not found in stocks in general.

Predicting factor returns fits into the broader literature on return predictability, which originated with the seminal paper by Campbell and Shiller (1988), who hypothesize that the dividend yield should have some preditive ability for the equity risk premium. The set of predictor variables has been extended throughout the years. Welch and Goyal (2008) provide a good overview of this set, which includes variables such as the short rate, term spread, stock variance, earnings/price ratio, book/market ratio, inflation and a measure of consumption. However, in the same paper Welch and Goyal (2008) raise serious doubts about the robustness of these findings. It seems as if in-sample returns are predictable, but this predictability vanishes when moving to OOS tests. Ang and Bekaert (2007) support these findings; only predictability by the short rate is found for short horizons.

A problem that may be simpler to solve is to estimate the direction of returns, since this is less noisy than actual returns. The simplest division of the direction of the market is to distinguish between a bull- and bear market, where the general momentum is up and down respectively. The research on the predictive ability of variables on market types is much less extensive than that on return prediction. Therefore the starting set of predictive variables is often taken as those arising from the literature on return predictability. Chen (2009) estimates multiple probit models with one macroeconomic predictor variable each to predict a bear market. He concludes that the macroeconomic variables indeed do a better job in predicting bear markets than in predicting returns. Nyberg (2013) also obtains promising results. He uses a more elaborate dynamic probit model that includes multiple macroeconomic and financial predictor variables to classify bull- and bear markets. He finds evidence of predictability both in-sample and OOS.

2.2 Markov-Switching Models in Return Predictability

In another strand of literature, market regimes are modeled as unobserved states, that are governed by a first order Markov-chain. In each of these states the distribution of the returns can have a different shape. These type of models are called Markov-switching models, and their first major appearance in econometrics is by Hamilton (1989), who uses it for estimation of the regime of the business cycle. Since then, the models have become quite popular in the literature on financial modelling, because they offer a way to parsimoniously capture some of the stylized behaviour of financial returns, such as fat tails, volatility clustering, skewness, and time-varying correlations (Ang and Timmermann, 2012).

Estimation of a Markov-switching model can be done without any predictor variables, because the probability of each state only depends on the state attained in the last period, the current observation, and the transition probabilities between the states. Such a Markov-switching model in its simplest form may already have an advantage over other models for predicting returns. Namely that it has comparable performance in-sample and OOS, at least for classifying the state of the market. Kole and Van Dijk (2017) find that the prediction of a bull- or bear market using a simple Markov-switching model OOS mostly is consistent with the full sample classifications for the returns of the S&P 500 index.

A Markov-switching model can however be made more complex by adding explanatory variables, which can enter the model in two ways. The first way is by entering as predictor variables in the equation for the regime-dependent return, with coefficients that are dependent on the state. This may alleviate the problem of time-varying relations between returns and predictor variables that Bender et al. (2018) identify as one of the main barricades for implementation of return prediction of factors. Schaller and Norden (1997) are the first to examine this for equity returns. They find a major discrepancy between the effect of the price/dividend ratio on returns in two market regimes.

The second way of using explanatory variables is by allowing the stochastic process of switching between the states to be endogenous. This boils down to endogenous transition probabilities between the states, which are dependent on the predictor variables. Diebold et al. (1994) are the first to propose this. Similar to the methods used in predicting the direction of the market mentioned earlier, the transition probabilities are estimated using a logit or probit specification. It however differs from these methods by using state dependent coefficients for the probabilities. Furthermore, the beliefs of the current market regime are incorporated.

Again, Schaller and Norden (1997) are the first to apply this method for equity returns. They use the price/dividend ratio to estimate the transition probabilities. It is found that allowing for time-varying probabilities leads to considerable absolute variation in the probability of staying in the ‘bad’ state and that there are relatively large peaks in the transition probability of moving to the ‘bad’ state preceding most recessions in the sample. Furthermore they pose the interesting question as to whether the price/dividend ratio has predictive ability only on the state, and not on returns (conditional on the state). It is however found that predictability is still present, when the price/ dividend ratio enters the transition probability equation.

Gray (1996) and Ang and Bekaert (2002b) incorporate time-varying transition probabilities with lagged values as predictors in a Markov-switching model for interest rates. Gray (1996) finds that doing so significantly improves the performance over the whole sample, using a likelihood ratio test to compare a model with varying probabilities to a model with constant probabilities. Ang and Bekaert (2002b) find that allowing transition probabilities to vary improves OOS forecasts, but only minimally improves regime classification. Kole and Van Dijk (2017) similarly find that using multiple predictor variables for the transition probabilities slightly improves OOS performance for predicting the S&P 500 index. A recurring finding in all studies incorporating time-varying transition probabilities is that most coefficients appearing in the probability equation have low significance, which Ang and Bekaert (2002b) impute to overparameterization. They further express concerns about estimation of models with time-varying probabilities in small samples. This concern calls for variable selection.

2.3 Markov-Switching Models in Asset Allocation

Since there is some success in predicting returns using Markov-switching models, a logical continuation is to try and use this knowledge in the context of asset allocation. It should therefore come as no surprise that there is ample literature on this subject. This branch of literature starts with a paper by Ang and Chen (2002), who find that correlations between U.S. stocks and the aggregate U.S. market are much greater for downside moves than for upside moves, considering a model with two regimes. The asymmetric movements are even more severe for small stocks, value stocks and past loser stocks. Ang and Bekaert (2002a) examine whether these dynamics can be exploited in an asset allocation problem. They develop a Markov-switching model for international asset returns and a risk-free asset. They find two market regimes: a high-volatility bear market and a bull market. They further conclude that the costs of ignoring the regimes are large when a risk-free asset can be held. In Ang and Bekaert (2004), their analysis is broadened by also investigating whether a Markov-switching model can be used for market timing. They find that a Markov-switching timing strategy outperforms a static strategy, because the model told the investor to switch to cash in a persistent high-volatility market.

While a bull- and bear market are the most recognisable regimes, the problem of Markov-switching models in asset allocation is not limited to two states. Guidolin and Timmermann (2007) find that four states – characterized as crash, slow growth, bull and recovery – are required to capture the joint distribution of stocks and bonds. They obtain the result that OOS there is economic importance in accounting for the presence of regimes in asset returns. They also examine the problem of asset allocation in a portfolio consisting of the size- and value factor in their 2008 paper. It is found that incorporating multiple regimes has a large impact on the optimal asset allocation, and the investor's utility. Again, four states are found to be optimal. OOS, portfolio strategies based on models that account for regimes dominate single-state benchmarks. Ammann and Verhofen (2006) study a similar problem, but use only two market regimes. Their model has weak forecasting abilities. However, using a strategy where 100% is invested in the value factor when a high-volatility regime is identified and 100% in the momentum factor when a low-volatility regime reigns, an equally weighted portfolio is outperformed.

3 Methodology

Multiple versions of a Markov-switching model will be estimated, to test for the optimal number of regimes, and to check whether and how explanatory variables influence estimates. To achieve this, the model will be differentiated in three ways.

First of all, a distinction will be made between $K = 1, 2, 3, 4$ regimes. The model with one regime serves as the benchmark, while the model with four regimes is hypothesized to produce the best fit, due to the results in earlier research on similar multivariate Markov-switching models. However, the simpler, and more intuitive division between two regimes – a bull and bear market – is also appealing. Also, a model with three regimes is considered, to accommodate for a transition regime. Secondly, models with constant and time-varying transition probabilities will be estimated. Thirdly, models with and without explanatory variables in the equation for the regime-dependent returns are considered.

In the remaining sections, the models will be denoted by their number of regimes (1,2,3,4), and the way that explanatory variables enter the model: no explanatory variables and constant transition probabilities (C); with explanatory variables for regime-dependent returns (CE) and with time-varying transition probabilities (TVTP).

3.1 The Structure of a Markov-Switching Model

The joint distribution of a vector of N returns (denoted in %), $\boldsymbol{r}_t = (r_{1,t}, \cdots, r_{N,t})'$ is modeled as a Markov-switching process driven by a state variable S_t with K states. It is important to note that this state variable is latent, and therefore also needs to be estimated. In this process, the vector of mean returns conditional on explanatory variables $\boldsymbol{\mu}_{k,t} = \mathbf{A}_k'\boldsymbol{x}_{t-1}$ and the covariance matrix $\boldsymbol{\Sigma}_k$ of the multivariate normally distributed return innovations $\boldsymbol{\varepsilon}_t = (\varepsilon_{1,t}, \cdots, \varepsilon_{N,t})' \sim \mathrm{N}(\mathbf{0}, \boldsymbol{\Sigma}_k)$ are allowed to differ between states:

$$\boldsymbol{r}_t = \mathbf{A}_k'\boldsymbol{x}_{t-1} + \boldsymbol{\varepsilon}_t \qquad k \in 1, \cdots, K \quad t \in 1, \dots, T. \tag{1}$$

The matrix $\mathbf{A}_k$ has n columns $(\alpha_{0i}, \alpha_{1n}, \cdots, \alpha_{Pn})'$ for $n = 1, \cdots, N$, with α_0 the intercept and $\alpha_1, \cdots, \alpha_P$ the coefficients for P explanatory variables. $\boldsymbol{x}_{t-1}$ is the column-vector $(1, x_{1,t-1}, \cdots, x_{P,t-1})'$, where only explanatory variables at time $t-1$ are used, such that the model can be used in 'real-time'. If $a_{pn} = 0$ for $p = 1, \cdots, P$ and $n = 1, \cdots, N$, the model collapses to the simpler model in which only the mean and covariance matrix are dependent on the regime.

The states follow a Markov chain governed by the transition probability matrix $\mathbf{P}_t$, with elements

$$p_{kj,t} = \Pr(S_t = k \mid S_{t-1} = j, \boldsymbol{x}_{t-1}) = \frac{e^{\boldsymbol{\beta}_{kj}\boldsymbol{x}_{t-1}}}{\sum_{i=1}^{K} e^{\boldsymbol{\beta}_{ij}\boldsymbol{x}_{t-1}}} \qquad k, j \in 1, \cdots, K \quad t \in 1, \ldots, T, \tag{2}$$

where $\boldsymbol{\beta}_{kj}$ is the vector $(\beta_{0,kj}, \beta_{1,kj}, \cdots, \beta_{P,kj})$, with $\beta_{0,kj}$ the regime-dependent intercept, and $\beta_{1,kj}, \cdots, \beta_{P,kj}$ the regime-dependent coefficients for P explanatory variables. For each departure state, the coefficients of one specific transition to state k are set to 0 for identification. This is the reference category. In this paper this is done for transitions to state one. Using the specification in equation 2, it is ensured that the columns of the matrix sum to one. If the restriction is imposed that $\beta_{p,kj} = 0 \quad \forall \quad k, j \in 1, \cdots, K \quad p \in 1, \cdots, P$, the model simplifies to the regular, and more widely used, model with a constant transition probability matrix.

The full model results in the following set of parameters to be estimated: $\boldsymbol{\theta} = \{\mathbf{A}_k, \boldsymbol{\Sigma}_k, \boldsymbol{\beta}_{kj} : k, j \in 1, \cdots, K\}$. This concerns a total of $K((N+1)P + n\frac{(N+1)}{2}) + K^2(1+P)$ parameters.

3.2 Parameter Estimation by the Expectation Maximization Algorithm

Because the problem under investigation suffers from incomplete data (namely the latent states), parameter estimation by regular maximum likelihood is computationally difficult, as the likelihood surface is often ill-behaved (Hamilton, 1989). The expectation-maximization (EM) by Dempster et al. (1977) however provides a solution for this problem. The algorithm consists of two steps that succeed each other iteratively. In the expectation (E) step, the expectation of the logarithm of the joint likelihood function of the observations and the states is taken, with respect to the states and given all the data. For this expectation, the values for parameters of the latest iteration are used. Next, in the maximization (M) step, this expectation is maximized with respect to the parameters of interest, providing updated parameter values. Hamilton (1990) shows that application of this algorithm for a Markov-switching model leads to convergence of all parameters, and more importantly, that this set of parameters leads to a maximum of the log-likelihood of the observations.

Usage of the EM-algorithm has multiple advantages. Firstly, standard properties of maximum likelihood estimators apply, because the log-likelihood is maximized. This is particularly useful for significance tests of parameters. Furthermore, it moves relatively quickly to decent values for the parameters, such that the choice of starting values is not that important. Lastly, the expected joint log-likelihood function has an attractive structure, which simplifies maximization. This will become clear soon.

In the following sections estimation of all parameters using the algorithm is discussed. By doing so, the particularities of the EM-algorithm applied to this case will become clear. Especially in the maximization step there are some extra steps compared to its first usage in estimating a Markov-switching model, due to implementation of time-varying transition probabilities with variable selection by the score test, and predictor variables for the mean with variable selection using the LASSO.

The algorithm can be initiated in both the expectation- and the maximization step. I choose to initiate with the expectation step, using randomly chosen starting values for the parameters. Convergence is attained when the relative change between the log-likelihood of two consecutive iterations is smaller than 10^{-8}.

3.2.1 Expectation Step (State Beliefs Parameters)

Using the set of estimates $\boldsymbol{\theta}^*$, from either initiation or the most recent maximization step, the expected value of the joint log-likelihood of the states and the observations is taken, with respect to the state and given all data. To arrive at this expectation, multiple recursions are needed that use expected values of the states given different information sets; these are the state beliefs parameters that are in our interest. These recursions – the Hamilton (1989) filter and Kim (1994) smoother – are discussed below, where I use notation and results based on the lecture notes by Lange (2020)

Let $\boldsymbol{S}_t$ be a K vector consisting of zeroes and a single one at position k^* indicating the true state $S_t = k^*$. Furthermore the expectation on the states at time t using information set $\mathcal{I}_a$ is denoted by the K vector $\hat{\boldsymbol{S}}_{t|a}$, with elements $\hat{s}^k_{t|a}, k = 1, \cdots, K$ the probability of being in state k. The Hamilton (1989) prediction step is:

$$\hat{\boldsymbol{S}}_{t+1|t} = \mathbf{P}_t \hat{\boldsymbol{S}}_{t|t}. \tag{3}$$

I define this entity as the predicted state. It is particularly useful if forecasting of returns is required, since it provides information about a future state.

Next, using information at time t, with the Hamilton (1989) updating step the beliefs can be updated to obtain the updated state:

$$\hat{\boldsymbol{S}}_{t|t} = \frac{\left(p\left(\boldsymbol{r}_t \mid S_t = 1\right), \cdots, p\left(\boldsymbol{r}_t \mid S_t = K\right)\right)' \odot \hat{\boldsymbol{S}}_{t|t-1}}{\left(p\left(\boldsymbol{r}_t \mid S_t = 1\right), \cdots, p\left(\boldsymbol{r}_t \mid S_t = K\right)\right) \hat{\boldsymbol{S}}_{t|t-1}}, \tag{4}$$

with $p\left(\boldsymbol{r}_t \mid S_t = k^*\right)$ the probability density of a normal distribution $\mathrm{N}(\boldsymbol{\mu}_{k,t}, \boldsymbol{\Sigma}_k)|_{k=k^*}$ evaluated at $\boldsymbol{r}_t$. This entity is of use if the best estimate of the current regime is needed, using all available information in real-time.

Lastly, in the Kim (1994) smoothing step, the expectations on the state at time t using all available information up to and including the last observation at time T can be derived as:

$$\hat{\boldsymbol{S}}_{t|T} = \hat{\boldsymbol{S}}_{t|t} \odot \mathbf{P}_t' \left(\hat{\boldsymbol{S}}_{t+1|T} \oslash \hat{\boldsymbol{S}}_{t+1|t} \right), \tag{5}$$

with $\mathbf{1}$ a column-vector of ones with length K. This expectation uses the most information and therefore is of interest when the best estimate of a state in the past is desired.

Next to the three types of state beliefs above, the matrix $\mathbf{P}_t^*$ with elements $p_{kj,t}^* = \Pr(S_t = k, S_{t-1} = j \mid \mathcal{I}_T)$ is derived:

$$\mathbf{P}_t^* = \mathbf{P}_t \odot \left(\hat{\boldsymbol{S}}_{t|T} \hat{\boldsymbol{S}}_{t-1|t-1}' \right) \oslash \left(\hat{\boldsymbol{S}}_{t|t-1} \mathbf{1}' \right). \tag{6}$$

It is only of use for calculating the expected value of the joint log-likelihood.

The joint log-likelihood of the observations and the states is given by:

$$\log f\left(\boldsymbol{r}_{1:T}, S_{1:T} \mid \mathcal{I}_T; \boldsymbol{\theta}; \boldsymbol{\rho}\right) = \sum_{t=1}^{T} \left[\sum_{k=1}^{K} \sum_{j=1}^{K} \delta_{kj,t} \left(\log\left[p\left(\boldsymbol{r}_t \mid S_t = k\right) \right] + \log\left[p_{kj,t} \right] \right) \right] + \sum_{k=1}^{K} \delta_{k,0} \log \rho_k, \tag{7}$$

where the subscript $1:T$ denotes all observations in the interval between 1 and T, $\delta_{kj,t}$ is an indicator that equals 1 when $S_t = k$ and $S_{t-1} = j$, and $\delta_{k,t}$ an indicator that equals 1 when $S_t = k$, and ρ_k is the probability that $S_0 = k$.

Using the just stated recursions, and latest set of parameters $\boldsymbol{\theta}^*$,the expected value of this joint log-likelihood given all data is:

$$\begin{aligned} \mathrm{E}_{S_t}[\ell\left(\boldsymbol{r}_{1:T}, S_{1:T} \mid \mathcal{I}_T; \boldsymbol{\theta}^*; \boldsymbol{\rho}\right) \mid \mathcal{I}_T] = & \sum_{t=1}^{T} \left[\sum_{k=1}^{K} \sum_{j=1}^{K} p_{kj,t}^* \left(\log\left[p\left(\boldsymbol{r}_t \mid S_t = k\right) \right] + \log\left[p_{kj,t} \right] \right) \right] \\ & + \sum_{k=1}^{K} \hat{s}_{0|T}^k \log \rho_k. \end{aligned} \tag{8}$$

After convergence of the EM algorithm, it follows that the log-likelihood of the observations:

$$\ell\left(\boldsymbol{r}_{1:T} \mid \mathcal{I}_{t-1}; \boldsymbol{\theta}\right) = \sum_{t=1}^{T} \log \left(\hat{\boldsymbol{S}}_{t|t-1}' \boldsymbol{p}\left(\boldsymbol{r}_t\right) \right), \tag{9}$$

where $\boldsymbol{p}\left(\boldsymbol{r}_t\right)$ is a vector with elements $p\left(\boldsymbol{r}_t \mid S_t = k\right)$ for $k = 1, \cdots, K$, is maximized.

3.2.2 Maximization Step

An attractive feature of the expected joint log-likelihood function in equation (8), is that it can be maximized for each set of parameters (those for the distribution of the returns, the transition probabilities, and the starting probabilities) separately. Solving the maximization problem for the transition probability and regime dependent distribution parameters will be discussed in detail in the next sections. The solutions for the starting states are also obtained. However, they are not discussed, because they have a minimal effect on the fit since a fairly long sample is considered.

3.2.2.1 Regime Dependent Distribution Parameters

To find the estimates of $\mathbf{A}_k$ and $\mathbf{\Sigma}_k$ for $k = 1, \cdots, K$, the following part of the expected joint log-likelihood is maximized with respect to the parameters:

$$\sum_{t=1}^{T} \sum_{k=1}^{K} \sum_{j=1}^{K} p^*_{kj,t} \left(\log \left[p\left(\boldsymbol{r}_t \mid S_t = k\right)\right]\right), \tag{10}$$

which, due to the property that $\sum_{j=1}^{K} p^*_{kj,t} = \hat{s}^k_{t|T}$, can be rewritten as:

$$\sum_{t=1}^{T} \sum_{k=1}^{K} \hat{s}^k_{t|T} \left(\log \left[p\left(\boldsymbol{r}_t \mid S_t = k\right)\right]\right). \tag{11}$$

Therefore, each state has the likelihood contribution

$$\sum_{t=1}^{T} \hat{s}^k_{t|T} \left(\log \left[p\left(\boldsymbol{r}_t \mid S_t = k\right)\right]\right), \tag{12}$$

which is exactly the same as the likelihood function of a weighted-least-squares (WLS) regression with weights $(\hat{s}^k_{1|T}, \cdots, \hat{s}^k_{T|T})'$. Consequently, for each state k the solution is the WLS solution:

$$\hat{\mathbf{A}}_k = (\mathbf{X}'\mathbf{D}(\hat{\boldsymbol{S}}^k_T)\mathbf{X})^{-1}\mathbf{X}'\mathbf{D}(\hat{\boldsymbol{S}}^k_T)\mathbf{R} \qquad \text{for } k = 1, \cdots, K, \tag{13}$$

$$\hat{\mathbf{\Sigma}}_k = (\mathbf{R} - \mathbf{X}\hat{\mathbf{A}}_k)'\mathbf{D}(\hat{\boldsymbol{S}}^k_T)(\mathbf{R} - \mathbf{X}\hat{\mathbf{A}}_k) \qquad \text{for } k = 1, \cdots, K, \tag{14}$$

with $\mathbf{X} = (\boldsymbol{x}_1, \cdots, \boldsymbol{x}_T)'$, where $\boldsymbol{x}_t = (1, x_{1,t}, \cdots, x_{P,t})$, $\mathbf{R} = (\boldsymbol{r}_1, \cdots, \boldsymbol{r}_T)'$, and $\mathbf{D}(\hat{\boldsymbol{S}}^k_T)$ a diagonal matrix with the vector $\hat{\boldsymbol{S}}^k_T = (\hat{s}^k_{1|T}, \cdots, \hat{s}^k_{T|T})'$ on the diagonal, the smoothed state probability of being in state k for each time t.

3.2.2.2 Transition Probability Parameters

The coefficients for the variables appearing in the equation for the transition probabilities are found by maximizing the following part of the expected joint loglikelihood:

$$\sum_{t=1}^{T} \sum_{k,j=1}^{K} p^*_{kj,t} \log\left(p_{kj,t}\right), \tag{15}$$

with $p_{kj,t}$ as in equation (2). To find a solution, a system of nonlinear equations needs to be solved and thus no analytical solution exists. Therefore this problem is solved numerically with the BFGS algorithm using `optim` in `R`.

The coefficients by themselves do not provide valuable inference on the effect of the variables on transition probabilities, because the model for the transition probabilities is non-linear. To retrieve this effect, the average marginal effects (AME) are calculated. The AME for variable x_p in the transition from state j to k can be interpreted as the average increase of the transition probability from state j to k caused by an unit increase of the variable of interest at the observed variables in the sample, ceteris paribus. It is calculated by taking the average of the partial derivative of equation (2), at each time t:

$$AME_{p,kj} = \frac{1}{T}\sum_{t=1}^{T}\frac{\partial p_{kj,t}}{\partial x_p} = \frac{1}{T}\sum_{t=1}^{T} p_{kj,t}\left(\beta_{p,kj} - \sum_{i=1}^{K} p_{ij,t}\beta_{p,ij}\right) \quad \text{for } k,j = 1,\cdots,K, \tag{16}$$

with $\beta_{p,kj}$ the coefficient for variable x_p in the transition from state j to k.

3.3 Computing Standard Errors

Since the EM-algorithm ensures maximization of the log-likelihood of the observations in equation (9), the standard properties of maximum likelihood estimators apply (Kole, 2019). Therefore, when the regularity conditions are satisfied, the estimated parameters $\hat{\boldsymbol{\theta}}$ converge in distribution in the following way:

$$\sqrt{T}\left(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0\right) \xrightarrow{d} \mathrm{N}\left(0, \boldsymbol{I}_0^{-1}\right), \tag{17}$$

with $\boldsymbol{\theta}_0$ the true parameters, and $\boldsymbol{I}_0$ the information matrix evaluated at the true parameters $\boldsymbol{\theta}_0$. This property is used to construct standard errors for the estimates, from the square root of the diagonal of the inverse of an estimator for the true information matrix. The standard errors are used to test whether the estimates of the parameters for the regime dependent mean, and those for the transition probabilities significantly differ from 0, using a t-test.

As an estimator for the information matrix, the outer-product-of-the-gradient (OPG) is used. This method provides an estimate for the information matrix. It is chosen over the other method, which is the negative of the Hessian, because estimation of the Hessian is computationally difficult in this case. The OPG estimate of the information matrix is:

$$\hat{\boldsymbol{I}}(\hat{\boldsymbol{\theta}}) = \frac{1}{T}\sum_{t=1}^{T}\boldsymbol{U}_t(\hat{\boldsymbol{\theta}})\boldsymbol{U}_{\mathrm{t}}(\hat{\boldsymbol{\theta}})', \tag{18}$$

where

$$\boldsymbol{U}_t(\boldsymbol{\theta}) = \frac{\partial \ell_t\left(\boldsymbol{r}_t \mid \mathcal{I}_{t-1};\boldsymbol{\theta}\right)}{\partial \boldsymbol{\theta}} \tag{19}$$

is the conditional score. The derivation of the conditional score can be found in appendix A.

3.4 Variable selection

As mentioned in section 3.1, the total of number of parameters that has to be estimated if the full set of explanatory variables is used in both the equation for the mean returns and transition probabilities is: $\theta = \{\mathbf{A}_k, \mathbf{\Sigma}_k, \boldsymbol{\beta}_{kj}; k, j \in 1,\cdots,K\}$. This concerns a total of $K((N+1)P + N\frac{(N+1)}{2}) +$

$K^2(1+P)$. Even though I only consider models with either time-varying transition probabilities or endogenous regime-dependent returns, and not both at the same time, this number can still become quite large if a model with three or four states is considered. Therefore, to avoid overparameterization and overfitting, decrease computation time, and to improve the interpretability of the model, variable selection procedures are used.

3.4.1 Regime Dependent Parameters

To apply variable selection for the regime dependent parameters, the grouped LASSO is used. Next to the advantages for variable selection named above, the LASSO is chosen because it both does continuous shrinkage and automatic variable selection simultaneously (Zou and Hastie, 2005). While this already is preferable over a discrete step-wise variable selection method (such as deployed for the transition probability parameters), due to high variability of these methods, this 'one-stop-shop' is particularly useful for the models used in this paper. Due to the relative computational slowness of the EM-algorithm, it is a big advantage to be able to obtain the set of selected variables all at once, instead of running the EM-algorithm until convergence time and again with different sets of parameters.

The grouped lasso is similar to the original LASSO, but instead of imposing a L1-norm on each parameter separately, it imposes an Euclidian norm (not squared) on groups of parameters (Hastie et al., 2017). This norm still has the ability to shrink parameters to exactly zero, but differs from the original LASSO in that it either shrinks all parameters of a group to exactly zero, or none. Here, this property is used to either include a regime-dependent variable for the multivariate return vector for all returns, or for none. Using the results from section 3.2.2.1, with the LASSO we maximize for each state $k \in 1, \cdots, K$:

$$\sum_{t=1}^{T} \hat{s}_{t|T}^{k} \left(\log\left[p\left(\boldsymbol{r}_t \mid S_t = k\right)\right]\right) + \lambda \sum_{j=1}^{p} \|\mathbf{A}_k^j\|, \tag{20}$$

where $\mathbf{A}_k^j$ is the j-th row of coefficient matrix $\mathbf{A}$ in state k and λ is the regularization parameter that determines the amount of regularization. Solving this weighted penalized likelihood problem can be done easily by the `R` package `glmnet`, by Friedman et al. (2010). This package deploys a cyclical coordinate descent algorithm, which is very efficient. Employment of the LASSO causes practically no extra computation time.

It remains to choose a value for the regularization parameter λ. Because a different LASSO model is fitted in each state, each state also has a different λ. For each state a sequence for λ is created that starts at the non-restrictive value of 0, and then moves in 9 equally spaced steps on the logarithmic scale to a value of λ that is very restrictive, such that no variables are selected.

This extreme value is found based on running the LASSO using values for the smoothed states provided by a model with no explanatory variables. Next, these λ's are grouped, such that there are 10 sets of λ's ranging from non-restrictive to very restrictive. Ideally, a grid-search over all these values would be conducted, but this results in a total of 10^K combinations, something that is computationally unfeasible. It is for the same reason that performance measures such as AIC and BIC – which are discussed in section 3.6 – are chosen over cross-validation to choose the final set of λ's.

The procedure is then to run the EM-algorithm one time for each set of λ's, using equation (20) instead of (12), and therefore with parameter estimates supplied by `glmnet` output instead of the analytical solution in equation (13). All other parameters are estimated in the usual manner. I choose the regularization parameter corresponding to the model that has the best values of the performance measures.

3.4.2 Transition Probability Parameters

For variable selection in the set of transition probability parameters $\boldsymbol{\beta}$, the score test is used. It is chosen over similar tests such as the Wald- or likelihood-ratio (LR) test, because the test only needs the restricted set of variables, instead of the unrestricted set in case of the Wald test, or both in case of the LR test. I furthermore use this method of variable selection over the LASSO, because the penalized version of the relevant part of the expected joint log-likelihood in equation (15) does not have a simple, equivalent, representation – such as the WLS interpretation for the regime-dependent distribution parameters – that is readily implemented in available software. This is due to the classes of the multinomial-logit problem being latent. Implementation of the LASSO for transition probabilities is beyond the scope of this research, but a very interesting direction for further research.

The score test tests the null hypothesis: $\boldsymbol{\beta} = \boldsymbol{\beta}_0$, where $\boldsymbol{\beta}_0$ is the original set of parameters with one variable that is considered for addition to the model restricted to be 0. Its statistic is chi-squared distributed with the number of restrictions –in this case one – as the degrees of freedom:

$$\boldsymbol{U}'(\hat{\boldsymbol{\beta}}_0)\boldsymbol{I}^{-1}(\hat{\boldsymbol{\beta}}_0)\boldsymbol{U}(\hat{\boldsymbol{\beta}}_0) \sim \chi^2_1, \tag{21}$$

where $\boldsymbol{U}(\hat{\boldsymbol{\beta}}_0)$ and $\boldsymbol{I}(\hat{\boldsymbol{\beta}}_0)$ are the score function and Fisher information matrix when the null hypothesis is true, evaluated at the restricted estimates (Breusch and Pagan, 1980). To estimate the Fisher information matrix, the OPG method is used. The derivation of the (conditional) score needed for the test can be found in appendix A

Using the score test, variable selection becomes a stepwise procedure in which the restricted model is estimated, after which the variable with the largest statistic from the score-test is added to the model. This procedure continues until no variable achieves a statistic that surpasses the 10% significance treshold. Because Davidson and MacKinnon (1983) find that the score test with the OPG estimate rejects the null too often when applied to small samples, and due to some numerical difficulties which will be elaborated on in section 5.3.1, as an intermediate step a LR test is conducted after the new model has been estimated, to check whether the increase in model fit is indeed significant.

3.5 Moving to Asset Allocation

For the asset allocation problem, I consider a buy-and-hold investor with mean-variance utility in a single- and multi-period setting. While there are more realistic utility functions and investment strategies, I choose this relatively simple specification because the purpose of this part of the paper is to show the power, or weakness, of a Markov-switching model of factor returns in an asset allocation setting, compared to the setting where regimes are unaccounted for. Furthermore, mean-variance preferences are still the benchmark model in the literature.

To implement the output of a Markov-switching model in a portfolio optimization problem, the estimates from the models for the regime-dependent moments such first have to be transformed to moments that are independent of the regimes. These moments will be derived in the next section, after which the portfolio optimization problem will be discussed.

3.5.1 Estimation of Moments

The moments of returns conditional on one-step-ahead-, current-time-, and smoothed state beliefs can be calculated using the law of iterated expectations. This results in the following equation for mean returns:

$$\hat{\boldsymbol{r}}_{t|a} = \mathrm{E}(\boldsymbol{r}_t \mid \mathcal{I}_a, \boldsymbol{x}_{t-1}) = [(\boldsymbol{r}_t + \varepsilon_t \mid S_t) \mid \mathcal{I}_a] = \sum_{k=1}^{K} \hat{s}^k_{t|a}(\boldsymbol{r}_t \mid S_t = k) = \sum_{k=1}^{K} \hat{s}^k_{t|a}\hat{\boldsymbol{\mu}}_{k,t} \quad \text{for } a \in t-1, t, T, \tag{22}$$

and variance:

$$\begin{aligned}
\hat{\boldsymbol{\Sigma}}_{t|a} = \mathrm{Var}(\boldsymbol{r}_t \mid \mathcal{I}_a) = \mathrm{E}(\boldsymbol{r}_t\boldsymbol{r}_t' \mid \mathcal{I}_a) - \mathrm{E}(\boldsymbol{r}_t \mid \mathcal{I}_a)[\mathrm{E}(\boldsymbol{r}_t \mid \mathcal{I}_a)]' = \mathrm{E}[\mathrm{E}(\boldsymbol{r}_t\boldsymbol{r}_t' \mid S_t) \mid \mathcal{I}_a] - \hat{\boldsymbol{r}}_{t|a}\hat{\boldsymbol{r}}_{t|a}' \\
&= \sum_{k=1}^{K}[\hat{s}_{t|a}^k(\mathrm{Var}(\boldsymbol{r}_t \mid S_t = k) + \hat{\boldsymbol{\mu}}_{k,t}\hat{\boldsymbol{\mu}}_{k,t}')] - \hat{\boldsymbol{r}}_{t|a}\hat{\boldsymbol{r}}_{t|a}' \\
&= \sum_{k=1}^{K}[\hat{s}_{t|a}^k(\mathrm{Var}(\hat{\boldsymbol{\mu}}_k + \boldsymbol{\varepsilon}_t \mid S_t = k) + \hat{\boldsymbol{\mu}}_{k,t}\hat{\boldsymbol{\mu}}_{k,t}')] - \hat{\boldsymbol{r}}_{t|a}\hat{\boldsymbol{r}}_{t|a}' \\
&= \sum_{k=1}^{K}[\hat{s}_{t|a}^k(\hat{\boldsymbol{\Sigma}}_k + \hat{\boldsymbol{\mu}}_{k,t}\hat{\boldsymbol{\mu}}_{k,t}')] - \hat{\boldsymbol{r}}_{t|a}\hat{\boldsymbol{r}}_{t|a}' \\
&\text{for } a \in t-1, t, T.
\end{aligned} \tag{23}$$

For predictions of returns over a longer horizon, from t up until $t+h$ denoted as $\boldsymbol{r}_{t:h}$, multi-period ahead predictions are required. Because predictions for $h > 1$ are dependent on future state beliefs, I use a Monte Carlo method to obtain these predictions. This method is only possible for methods without explanatory variables, because these should also be predicted. In this paper, only these models are considered for multi-period predictions.

The Monte Carlo algorithm for estimating h period ahead returns consists of the steps as in algorithm 1 in appendix C.

3.5.2 The Portfolio Optimization Problem

I consider a buy-and hold investor with mean-variance utility who can allocate his/ her wealth across a 'risk-free' asset – a one month treasury bill – with h period return $r_{t:h}^f$ and the factors, with h period return $\boldsymbol{r}_{t:h}$. The considered horizons are 1, 4, 13, 26, and 52 weeks, corresponding to one week, one month, quarter of a year, half a year and full year returns, respectively. In the case of one week ahead returns, I use the notation: $\boldsymbol{r}_{t:1} = \boldsymbol{r}_{t+1}$.

(S)he allocates wealth over the factors using the weights $\boldsymbol{w}_t$. A no-short-selling and no-borrowing constraint is imposed, such that $0 \le w_{i,t} \le 1$ for $w_{i,t} \in \boldsymbol{w}_t$. This is done to avoid extreme weights, and thus to obtain more neat results.

The investor solves the following maximization problem of his/ her expected utility at time t:

$$
\begin{aligned}
\max_{\boldsymbol{w}_t} \quad \mathrm{E}\left[U(r^p_{t:h}) \mid \mathcal{I}_t\right] = \quad & \mathrm{E}\left[r^p_{t:h} \mid \mathcal{I}_t\right] - \frac{1}{2}\gamma \mathrm{Var}\left[r^p_{t:h} \mid \mathcal{I}_t\right] \\
= \quad & \mathrm{E}\Big[(1 - w_{MKT})r^f_{t:h} + w^{MKT}_t r^{MKT}_{t:h} + \boldsymbol{w}^{*\prime}_t \boldsymbol{r}^*_{t:h|t} \mid \mathcal{I}_t\Big] \\
& - \frac{1}{2}\gamma (w^{MKT}_t, \boldsymbol{w}^{*\prime}_t) \mathrm{Var}\left[r^p_{t:h} \mid \mathcal{I}_t\right] (w^{MKT}_t, \boldsymbol{w}^{*\prime}_t)' \\
= \quad & (1 - \boldsymbol{w}'_t \mathbf{1}) r^f_{t:h} + \boldsymbol{w}'_t (\hat{r}^{MKT}_{t:h|t}, \hat{\boldsymbol{r}}^{*\prime}_{t:h|t} + r^f_{t:h}) - \frac{1}{2}\gamma \boldsymbol{w}'_t \hat{\boldsymbol{\Sigma}}_{t:h|t} \boldsymbol{w}_t \\
= \quad & r^f_{t:h} + \boldsymbol{w}'_t (\hat{r}^{(e)MKT}_{t:h|t}, \hat{\boldsymbol{r}}^{*\prime}_{t:h|t}) - \frac{1}{2}\gamma \boldsymbol{w}'_t \hat{\boldsymbol{\Sigma}}_{t:h|t} \boldsymbol{w}_t, \\
\text{s.t.} \quad & 0 \leq w_{i,t} \leq 1 \quad \text{for} \quad w_{i,t} \in \boldsymbol{w}_t,
\end{aligned}
\tag{24}
$$

with $r^p_{t:h}$ the h-period portfolio return, γ the coefficient of risk aversion, $\boldsymbol{w}^*$ and $\boldsymbol{r}^*$ vectors with the weights and returns for portfolios (HML, QMJ, SMB, UMD) respectively, and superscript $^{(e)}$ denoting excess returns. The second equality arises because all portfolios except the market portfolio are zero-cost portfolios. Thus, any weight assigned to one of these portfolios does not impact the budget constraint. In the next two lines this equation is rewritten in such a way that it is in the form of a quadratic programming problem. This problem is solved using the `R` package `quadprog`. A moderate level of risk aversion is considered, with $\gamma = 3$.

3.6 Performance Measures

The performance measures that are used fall into two categories that serve different purposes. First of all, because interest lies in the regime-dependent factor dynamics, I want to obtain a model that provides the most accurate presentation of the true process, while remaining relatively parsimonious, such that the model remains interpretable. The first class of performance measures should evaluate this model fit. Secondly, as the objective is to use the models in an portfolio optimization problem, the second class of performance measures should compare and measure the performance of the portfolios formed with different types of models.

3.6.1 Model Specification

As a general indication of model fit, the log-likelihood will be used, where a higher loglikelihood indicates a better fit. A drawback of this simple measure is that adding variables will always increase the loglikelihood. Therefore information criteria will also be used to evaluate the model, as they trade-off the loglikelihood with model parsimony. The information criteria have the form:

$$
IC = \alpha - 2\ln(\hat{L}), \tag{25}
$$

where α is some constant, and $\hat{L}$ is the likelihood of an estimated model. The Akaike- and Bayesian Information criterion (AIC & BIC respectively) are used. AIC uses $\alpha = 2|\boldsymbol{\theta}|$ while BIC uses $\alpha = |\boldsymbol{\theta}| \ln T$, where $|\cdot|$ denotes the cardinality of a set. Generally AIC selects the model that has the best predictive fit, while BIC searches for the true model. In practice, this implies that BIC tends to select model that is more parsimonious. A lower score on either one of these information criteria indicates a better fit.

As a further assessment of the best specification of a Markov-switching model, the regime classification measure (RCM) is used, which is a measure designed specifically to test for the optimal number of regimes in a Markov-switching model. It can be defined in multiple ways, but the definition by Guidolin (2009) is chosen, since it is best suited for models with multiple regimes. It is defined as:

$$RCM(K) = 100 \left\{ 1 - \frac{K^{2K}}{(K-1)^2} \frac{1}{T} \sum_{t=1}^{T} \prod_{k=1}^{K} \left[\hat{s}_{t|T}^{k} - \frac{1}{K} \right]^2 \right\}. \tag{26}$$

The intuition behind it is that the majority of smoothed probabilities will be close to 0 and only one close to 1 at each time t, when the model is able to distinguish very sharply between states. In the exact case, this will lead to a RCM score of 0. On the other hand, when the model provides the most naive smoothed probability of being in a state, namely $\frac{1}{K}$, a RCM score of 100 is obtained. Thus, a lower score is preferred.

Lastly, I use a density specification test to test whether the null hypothesis that the density of a model is correctly specified. This is an important measure, for one because it provides the most information about the correctness of the regime dynamics. Furthermore, correct specification of the density is of major importance in the portfolio optimization problem. Because a mean-variance investor is considered, not only the mean, but also the second moment of the distribution is of interest. A density specification test thus delivers a more reasonable way of assessing the model performance than a performance measure that solely focuses on a point forecast of the mean, such as RMSE.

Guidolin and Ono (2006) propose a form of the density specification test for Markov-switching models. It can be used for both the univariate and multivariate case. While the multivariate test is most useful in evaluating the model fit, the univariate test can provide more information about which variable is responsible for a possible rejection. For the univariate test each entry in the vector of predicted returns $\hat{\boldsymbol{r}}_{t+1|t} = (r_{1,t+1|t}, \cdots, r_{N,t+1|t})$ is transformed to a variable $z_{n,t+1}$ for

$n \in 1, \cdots, N$. It is calculated as:

$$\begin{aligned} z_{n,t+1} = \Pr\left(r_{n,t+1} \leq \hat{r}_{n,t+1} \mid \mathcal{I}_t\right) &= \sum_{k=1}^{K} \Pr\left(r_{n,t+1} \leq \hat{r}_{n,t+1} \mid \mathcal{I}_t, S_{t+1} = k\right) \Pr\left(S_{t+1} = k \mid \mathcal{I}_t\right) \\ &= \sum_{k=1}^{K} \Phi\left(\hat{\sigma}_{k,n}^{-1}\left[r_{n,t+1} - \hat{\mu}_{n,k,t+1}\right]\right) \hat{\boldsymbol{S}}_{t+1|t}, \end{aligned} \tag{27}$$

where Φ is the univariate standard-normal cumulative distribution function (cdf), and $\sigma_{k,n}$ is the estimate for the volatility of return n in state k, taken as the square root of the n-th entry on the diagonal of the estimate for the covariance matrix in state k; $\hat{\boldsymbol{\Sigma}}_k$. Predicted states are used here, since these are the states that can be used in a real-time OOS experiment.

This variable is uniformly independently and identically distributed (IID) on the interval [0,1], under the null of correct specification of the model. It has the intuition that deviations between realized and predicted returns should be conditionally normal. Therefore, they should describe a uniform distribution once they are 'filtered through' an appropriate Gaussian cdf. The z_t should be IID because when the model is correctly specified, the errors should have no structure and should be unpredictable (Guidolin and Ono, 2006).

Next, all $z_{n,t}$ are transformed back to a variable $z^*_{n,t}$,using $z^*_{n,t} = \Phi^{-1}(z_{n,t})$. Under the null, it should be IID standard normal. On this variable, a Shapiro-Wilk (1965) test for normality and IID-ness is conducted.

For the multivariate test, I deviate from Guidolin and Ono (2006) and use the transformation:

$$z_{t+1} = \sum_{k=1}^{K} F_{\chi N}\left[\left(\boldsymbol{\Omega}_k^{-1}\left[\boldsymbol{r}_{t+1} - \hat{\boldsymbol{\mu}}_{k,t+1}\right]\right)' \left(\boldsymbol{\Omega}_k^{-1}\left[\boldsymbol{r}_{t+1} - \hat{\boldsymbol{\mu}}_{k,t+1}\right]\right)\right] \hat{\boldsymbol{S}}_{t+1|t}, \tag{28}$$

where $F_{\chi N}$ is the cdf of a chi-squared distribution with N degrees of freedom, and $\boldsymbol{\Omega}_k$ is the Choleski factorization of the estimate for the regime dependent covariance matrix $\hat{\boldsymbol{\Sigma}}_k$. It is used that under the null of correct specification, after substraction of the mean and transformation by the inverse of the Choleski factorization, the deviations should be multivariate standard normal. Since multivariate standard normal variables are mutually independently univariate standard normally distributed, the sum of the N squares should have a chi-squared distribution with N degrees of freedom, and z_t therefore again is uniformly distributed on [0,1] under the null. Subsequently, the same steps are conducted as in the univariate case, to arrive at a single statistic.

This rather cumbersome transformation is preferred over the multivariate specification proposed by Guidolin and Ono (2006), who use a multivariate normal cdf in a similar way as in the univariate case. However, it seems that due to the high multivariate dimension of the multivariate normal distribution, after the transformation the values for z_t are generally too low. This suspicion is confirmed by a quick experiment using simulated data from a multivariate normal distribution;

using the test multivariate normality is rejected. More research on this phenomenon is required.

Lastly, even though interest is in the correct estimation of the entire density, the RMSE is still used to measure the accuracy of point estimates. It is given by:

$$RMSE_n = \sqrt{\frac{\sum_{t=1}^{T}(\hat{r}_{n,t} - r_{n,t})^2}{T}} \tag{29}$$

3.6.2 Portfolio Optimization Performance Measures

Since an investor with mean-variance preferences is considered, natural performance measures for the portfolio optimization problem are the mean, variance, Sharpe-ratio (SR) and of course the realized utility. The SR is the excess portfolio return divided by the standard deviation of the portfolio returns:

$$SR = \frac{r^{p,(e)}}{\sigma^p}. \tag{30}$$

The investor's realized utility at time t is given by:

$$U(r^p_{t:h}) = \quad r^p_{t:h} - \frac{1}{2}\gamma \boldsymbol{w}_t'(\boldsymbol{r}_{t:h} - \hat{\boldsymbol{r}}_{t:h})(\boldsymbol{r}_{t:h|t} - \hat{\boldsymbol{r}}_{t:h|t})'\boldsymbol{w}_t, \tag{31}$$

with $r^p_{t:h} = r^f_{t:h} + \boldsymbol{w}_t'(r^{(e)MKT}_{t:h|t}, \boldsymbol{r}^{*\prime}_{t:h|t})$. The average value over all t is taken as a measure of performance, with all returns in decimal rates.

While all the above-mentioned performance measures will be reported, I choose to only test whether there is a significant difference between the Sharpe-ratios of different models. This is because the SR incorporates both measures important to a mean-variance investor, and therefore is a complete measure. Only the utility incorporates more information, namely the coefficient of risk aversion. But, since only one value of this coefficient is used, testing significance of differences between Sharpe-ratios is deemed sufficient. To test the difference, the test by Ledoit and Wolf (2008) is used, since it can deal with the heavy-tailed and time-series nature of returns.

3.7 Out-of-Sample Procedure

To obtain an OOS fit of the models, an expanding window method is used. The models are trained using all data preceding the first week of July 2003, which constitutes half of the sample. Next, 10 equally sized succeeding divisions in the remaining sample are made. The models are trained on the latest vintage, and tested on the following division of the data. This course division of the data is chosen due to computational constraints. Running the EM-algorithm for as many times as half the sample size would take very long, especially for the models which contain an element of variable selection.

4 Data

4.1 Dependent Variables

The data for the factor returns is collected from the dataset that is related to the paper "Quality Minus Junk" by Asness et al. (2019).[1]. It contains daily factor returns starting on the first of January 1926 and ending in the last week of December 2020 for the value, momentum, quality and market factor. This dataset is supplemented with daily data on the size factor from Kenneth French's website [2], because for some reason, the series for the size factor was incomplete. Due to availability of data on the explanatory variables, only data starting on the fourth of January 1986 is used. Furthermore, the data is transformed to weekly returns, to reduce the noise. This results in a total of 1,828 observations for each series. I argue that starting observations at this later date is not detrimental for the research, because factor premia were not discovered before the seventies, and the period spanning the dataset still includes much variation in the factor returns.

The value, momentum, and quality factor returns are constructed using value-weighted portfolios that are rebalanced monthly, sorted on size, book-to-market ratio, prior 12 month returns (skipping the most recent month), and an average of four aspects of quality respectively. The four aspects of quality are: profitability; growth; safety and payout. For the measures of these aspects, and the precise construction of the quality score, I refer to the original article by Asness et al. (2019). All sorts are conditional, which is preferred by the authors because it ensures a balanced number of securities in each portfolio. The firms are first sorted into two size portfolios, of which the breakpoint is the median NYSE market equity. Next, they are sorted into three portfolios using the second variable that depends on the factor of interest. For factor X, this results in returns of six portfolios: Small, Low X; Small, Neutral X; Small, High X; Big, Low X; Big, Neutral X; Big, High X. Using these portfolios, long-short portfolios are created, which will henceforth be called quality minus junk (QMJ), high minus low (HML), and up minus down (UMD) for the quality, value, and momentum factors respectively.

The factor return r_X, for factor $X \in \{\text{HML}, \text{UMD}, \text{QMJ}\}$ is constructed as:

$$r_X = \frac{1}{2}(\text{ Small, High X } + \text{ Big, High X }) - \frac{1}{2}(\text{ Small, Low X } + \text{ Big, low X }). \tag{32}$$

In addition, the size (SMB) factor is the average return of the small portfolios minus the average

[1] https://www.aqr.com/Insights/Datasets/Quality-Minus-Junk-Factors-Daily

[2] https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

Table 1: The historical mean, covariance, and correlation matrix of all factors over the full sample.

	QMJ	SMB	HML	UMD	MKT
Mean	0.093	0.001	0.003	0.123	0.167

Covariance & correlation matrix	QMJ	SMB	HML	UMD	MKT
QMJ	1.28	−0.34	−0.17	0.31	−0.52
SMB	−0.52	1.81	−0.11	−0.01	0.11
HML	−0.23	−0.18	1.52	−0.33	−0.12
UMD	0.71	−0.02	−0.84	4.11	−0.12
MKT	−1.39	0.36	−0.35	−0.57	5.61

Note: The colored cells display the correlations. A darker red implies a larger positive correlation, while a darker blue implies a larger negative correlation.

3

return of the big portfolios that are formed on size and book-to-market value:

$$\begin{aligned} r_{SMB} = \; & \frac{1}{3}(\text{ Small Value } + \text{ Small Neutral } + \text{ Small Growth}) \\ & -\frac{1}{3}(\text{ Big Value } + \text{Big Neutral } + \text{ Big Growth}). \end{aligned} \tag{33}$$

Lastly, the market factor return r_{MKT} is the value weighted return of all available stocks, minus the one month treasury-bill rate.

The mean and volatility, together with the covariances and correlations of each of the factors can be found in Table 1. These moments will later on be used as estimates for a model that does not acknowledge regimes. It can be seen that UMD and MKT achieve the highest mean returns, which are accompanied by the highest volatilities.

4.2 Explanatory Variables

The explanatory variables that are considered are a mixture of variables that are identified as possible predictors from literature on predicting market regimes, factor timing, and return predictability. They can be divided into measures of five broad categories: financial conditions, economic conditions, market sentiment, valuation, and momentum. All data is collected with daily frequency. Again, it is transformed to weekly data. The final dataset ranges from the first week of January 1986 up to and including the last week of December 2020, and is assumed to be known with a lag of one week. Thus, at time t, information of time $t-1$ is known.

To measure financial conditions, the TED spread is used, which is the difference between the 3-month treasury bill yield and the 3-month LIBOR yield. Because it measures the difference between the 'risk-free' treasury rate and the rate for interbank loans, it is a measure of credit risk. For economic conditions, the yield spread is used, which past research has identified as a leading indicator for the economic cycle. It is the yield spread between 3-month treasury bills and 10-year treasury notes. As longer-term debt should carry more risk, generally speaking they should have a higher rate. Therefore, a tighter spread can be an indicator of a short-term economic downfall. Data on both these series is obtained from the FRED database.[4] They are transformed from daily to weekly data by taking the weekly average. The measure used for market sentiment is the VXO, which is a forward-looking measure of market volatility based on the implied volatility of S&P 100 stocks. It is similar to the wider known VIX, but it has available data for a longer period, and therefore is preferred. Data of the VXO is obtained from the CBOE, which constructs the index.[5] Again, the weekly average is taken from the daily data. As a measure of valuation, ideally a metric would be used that indicates the valuation of a factor relative to other factors, such as the dividend yield, book-to-price ratio, or price-to-earnings ratio of all stocks in the factor portfolios. However, since I have no information on the individual stocks in the portfolios, and the factor portfolios are long-short portfolios, these measures are not possible to implement. Therefore I measure the valuation of a factor relative to its own historical value. This is done by taking the relative difference of the inflation adjusted value of a factor and its 2500-day moving average, starting with a base level of 100 at the start of the series in 1926. The last value of the week is taken. These variables are denoted as *MA 2500 ...* . Lastly, the one-month return of each factor is taken as a measure of momentum, because recent research by Gupta and Kelly (2019) finds that the effects of factor momentum are strongest over a one-month look back period. Again, the last value of the week is taken. I call these variables *prior r ...* .

Descriptive statistics for the explanatory variables can be found in appendix B.

5 Results

5.1 Model Specification

I base the best model specification on three types of criteria: (i) statistical measures, (ii) economic significance, (iii) economic interpretation. For statistical measures, information criteria, the RCM and a density specification test are used. To evaluate economic significance, in-sample measures of portfolio performance are used. While these measures are not relevant for investors – who are only interested in the performance OOS– they are useful for model selection. For the economic interpretation, an in-depth look of model estimates is needed. The first two criteria are discussed

[4]The TED spread: `https://fred.stlouisfed.org/series/TEDRATE`. The yield spread: `https://fred.stlouisfed.org/series/T10Y3M`

[5]`https://www.cboe.com/tradable_products/vix/vix_historical_data/`

Table 2: Performance measures for a Markov-switching model with constant transition probabilities and no explanatory variables with four, three, and two states (4C, 3C, 2C respectively), a model with three states and explanatory variables for the regime-dependent mean (3CE), and a model with time-varying transition probabilities (3TVTP), estimated on the full sample.

	4C	3C	2C	1C	3CE	3TVTP
			Model evaluation measures			
Number of parameters	96	69	44	20	139	90
Log-likelihood.	-14060	-14332	-14675	-16266	-14253	-14241
RCM	23.54	23.17	16.33	100	22.65	30.89
AIC	28312	28802	29438	32577	28784	28661
BIC	28841	29182	29680	32698	29549	29157
P-value Dens. Spec. Test ($\times$ 100)	0.0061	0.0516	$3.38 \cdot 10^{-7}$	$3.57 \cdot 10^{-29}$	0.0057	0.4459
			Portfolio optimization performance measures			
Mean excess portfolio return (%)	0.118	0.112	0.108	0.087	0.181	0.108
Volatility (%)	0.524	0.507	0.493	0.49	0.614	0.492
Realized utility ($\times$ 10,000)	17.401	16.777	16.409	14.357	23.499	16.439
Sharpe ratio (annualized)	1.627	1.588	1.576	1.285	2.121	1.585
			Sharpe ratio difference			
3TVTP	0.042	0.003	-0.009	-0.3^{**}	0.536^{***}	
3CE	-0.494^{***}	-0.533^{***}	-0.545^{***}	-0.836^{***}		
1C	0.342^{**}	0.303^{**}	0.291^{**}			
2C	0.051	0.012				
3C	0.039					

Note: The first block shows model evaluation performance measures. The second block shows portfolio performance measures for a portfolio formed as in equation (24), with state independent moments calculated using predicted state beliefs $\hat{\boldsymbol{S}}_{t+1|t}$ as in equations (22, 23). The Sharpe-ratio is annualized by multiplying by $\sqrt{52}$. The third block shows the differences in Sharpe ratios, with significance levels of a Ledoit and Wolf (2008) test. $^{***}, ^{**}, ^{*}$ display significance at the 10%, 5%, and 1% level respectively.

here, with their measures found in Table 2, while the economic interpretation is addressed in the next section.

In the first block of Table 2, it can be seen that the information criteria unambiguously point to a model with four states; the 4C model has both the lowest AIC and BIC. The decreases in information criteria are more extreme when moving from two states to three states, compared to moving from three to four states. The RCM points out that a model with two regimes distinguishes most sharply between the states, while the measure is approximately the same for the four-state and three-state model. For all models this measure is however relatively low; in the article by Guidolin (2009) where the RCM as in equation (26) is proposed, it spans a much wider range. Testing multiple specifications of a Markov-switching ARCH model with two, three, or four states on the bivariate distribution of U.S. bond and equity returns, the RCM varies from 20 up to 70, with averages around 40.

The multivariate density specification test however favors a model with three states, albeit

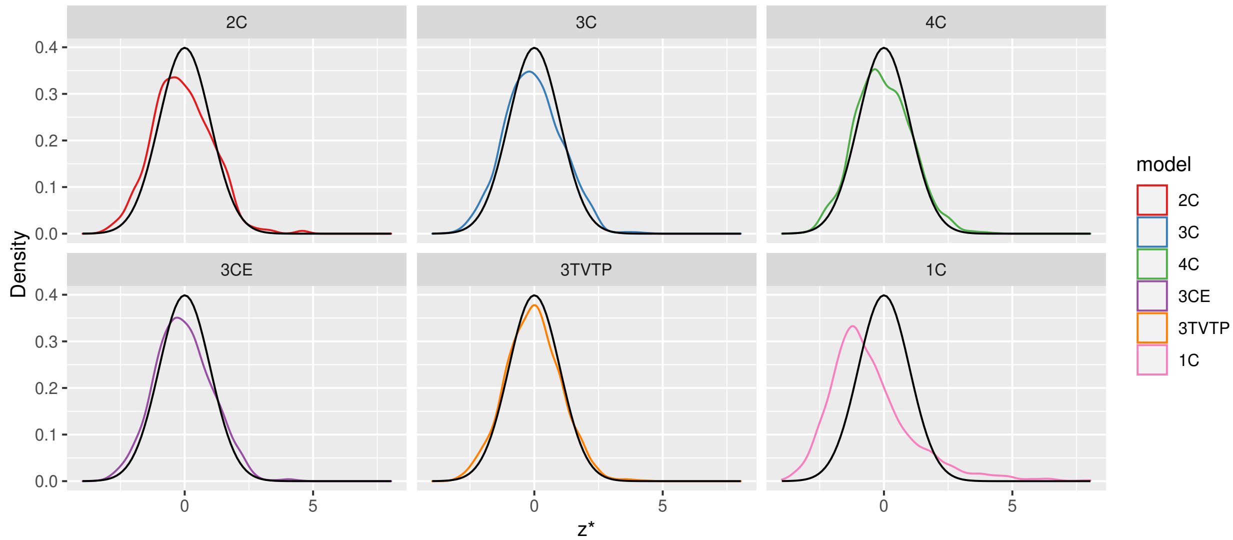

Figure 1: The density of z^*, as described in section 3.6.1, for all considered models, together with the standard normal density.

the null of correct specification is rejected for all models. This rejection however is substantially stronger for a model with two states than for a model with three or four states, with a p-value that is a factor 10,000 smaller than that of the four state model. The three state model has a p-value that is approximately a factor 10 larger than that of the four state model.

A visual representation of the misspecifications of the models can be seen in the upper row of Figure 1, where the density of the transformed z^* of the simple Markov-switching models (2C, 3C, 4C) can be found. Indeed, the density of z^* from the 3C resembles the standard normal distribution most closely, while the 2C model is most off. Either way, the density of z^* of the Markov-switching models resemble the standard normal density much more closely than a model that does not take regimes into account (1C), which can be seen by the highly skewed distribution in the bottom-right of the figure.

The source of the misspecifications is further examined by considering an univariate density

Table 3: The p-values for the univariate density specification test as in equation (27), performed on each of the portfolios of the different models.

	p-value ($\times$ 100)				
Model	QMJ	SMB	HML	UMD	MKT
4C	0.4802	0.0238	15.1460	0.0005	0.0000
3C	0.8565	0.0197	10.0412	0.0002	0.0001
2C	0.5253	0.0000	12.2928	0.0000	0.0000
3TVTP	24.1409	0.4201	7.1809	0.0011	0.0000
3CE	0.9708	0.0085	27.4540	0.0015	0.0001

specification test for each series of returns, of which the results can be found in Table 3. Generally, for all models the most evidence against misspecification is found along the dimensions of the UMD and MKT portfolio. I only fail to reject the null of correct specification for the HML portfolio in all models.

While all performance measures on model fit greatly improve when multiple regimes are considered in comparison to the 1C model, this is not translated into improvements in point prediction of returns. As can be seen in Table 4, the RMSE of all Markov-switching models are similar to the RMSE of the 1C model. To be clear, this is a model which just takes the historical mean and covariance matrix as in Table 1 as its inputs.

But, just as with model fit, acknowledging regimes does lead to significant improvements in economic value, when applied in in-sample portfolio optimization. As can be see in the second block of Table 2, all models produce portfolios with significantly higher Sharpe ratios than a portfolio produced by estimates from the 1C model. There are however no notable differences between the added economic value of a model with two, three, or four states, in-sample. All models achieve similar Sharpe ratios and realized utilities.

The superior economic values that Markov-switching models have emphasizes the importance of the role that both the first and second moment play in mean-variance optimization. It is the interplay between them that establishes the portfolio weights; while point prediction of mean returns is poor, the ability of Markov-switching models to vary the covariance matrix produces weights that allow them to reach higher realized portfolio returns, while keeping the volatility equal. This has higher Sharpe ratios as a result.

Given the discussion in the previous paragraphs, the specification with three states is preferred, and is used for further research in this paper. For one, it has the least evidence against not being specified correctly, and from the densities in Figure 1 it is seen that the z^* from the 3C model resemble the normal distribution most closely. Furthermore, in comparison to a four-state model it offers a more parsimonious way of modelling a process that is very similar in a large part of the sample, as will become clear in the next section. Besides, the choice of three over four regimes does

Table 4: The RMSE for each factor in different models.

	RMSE				
Model	QMJ	SMB	HML	UMD	MKT
4C	1.13	1.34	1.23	2.02	2.37
3C	1.13	1.35	1.23	2.02	2.37
2C	1.13	1.34	1.23	2.02	2.37
1C	1.13	1.35	1.23	2.03	2.37
3CE	1.11	1.33	1.22	2.00	2.36
3TVTP	1.14	1.34	1.24	2.02	2.37

not affect the portfolio performance significantly. Parsimony is an important argument here; the multitude of parameters in a four state model causes difficulty with estimation of the parameters, which slows convergence of the EM algorithm. This is particularly problematic when time-varying transition probabilities are considered, because numerical optimization is involved here.

5.2 Model Estimates

As a first step in economic interpretation of the regimes, I will classify them based on their intra- and inter-regime dynamics, which are described in Table 5 and Figure 2. In the table the regime-dependent parameters and transition matrix for the preferred specification with three states – along with those of the specifications with two or four states for comparison – can be found, while Figure 2 shows the cumulative returns of each of the factors together with the sequence of regimes with the highest smoothed probability for each model.

Based on the the diagonal of the regime-dependent covariance matrices of the 3C model in the middle column of the table, it can be seen that for the three-state model the identified regimes are most easily distinguished by their volatility. In regime one, volatility is low, and it increases with each regime. Furthermore, the sequence of most likely smoothed states of the 3C model in Figure 2 displays an interesting pattern: regime two consistently is an intermediate for a transition between regime one and three. What becomes apparent from the figure is that either one of these regimes transitions to regime two, after which the transition is either back to the 'old' regime, or to a 'new' state. Regime-switches very seldom happen directly between these regimes. The need of a transitory instead of sudden change in volatility is in line with the stylized fact of volatility having a long memory.

Due to the differences in volatility, and the structure of the most likely prevailing regimes, I choose to label the regimes in the following manner: *low-volatility* regime one, transitional *medium-volatility* regime two, and *high-volatility* regime three. Further characteristics and economic interpretations of the regimes will be elaborated on in the following paragraphs.

In *low-volatility* regime one, the transition probability matrix in the upper block of Table 5 shows that persistence is very high, with a 97.6% chance of remaining in the same regime. The transition probabilities to other regimes are small, although transitioning to the medium-volatility regime is twice as likely as transitioning to the high-volatility regime. Regime one most closely resembles the classic 'bull' market: it has significantly positive mean returns for MKT, just as for the other factors except SMB. Figure 2 shows that regime one coincides with periods of steady, low volatility, growth for most factors, such as a large part of the period between the start of the sample and 1999, and the period starting in the second half of 2003 and ending after the first half

Table 5: The transition matrices and moments of the 2C, 3C, and 4C models, estimated on the full sample.

Transition matrix

4C

to \ *from*	Regime 1	Regime 2	Regime 3	Regime 4
Regime 1	0.98	0.004	0.009	0.005
Regime 2	0.007	0.893	0.187	0.015
Regime 3	0.003	0.086	0.808	0.003
Regime 4	0.01	0.000	0.012	0.978

3C

to \ *from*	Regime 1	Regime 2	Regime 3
Regime 1	0.976	0.029	0.026
Regime 2	0.015	0.894	0.159
Regime 3	0.009	0.08	0.811

2C

to \ *from*	Regime 1	Regime 3
Regime 1	0.949	0.139
Regime 3	0.051	0.861

Mean

4C

	QMJ	SMB	HML	UMD	MKT
Regime 1	-0.55	0.080	0.107***	0.113**	0.235***
	(0.041)	(0.051)	(0.039)	(0.056)	(0.079)
Regime 2	0.125***	-0.02	-0.181***	0.389***	0.18**
	(0.044)	(0.054)	(0.045)	(0.078)	(0.093)
Regime 3	0.177	0.099	0.198	-0.658**	-0.205
	(0.150)	(0.171)	(0.162)	(0.294)	(0.312)
Regime 4	0.153***	-0.97**	0.009	0.238***	0.277***
	(0.024)	(0.046)	(0.036)	(0.036)	(0.077)

3C

	QMJ	SMB	HML	UMD	MKT
Regime 1	0.044**	0.002	0.065***	0.183***	0.283***
	(0.022)	(0.034)	(0.025)	(0.031)	(0.049)
Regime 2	0.136***	-0.05	-0.198***	0.407***	0.195***
	(0.044)	(0.054)	(0.045)	(0.079)	(0.099)
Regime 3	0.183	0.099	0.185	-0.655**	-0.299
	(0.142)	(0.161)	(0.155)	(0.276)	(0.294)

2C

	QMJ	SMB	HML	UMD	MKT
Regime 1	0.077***	-0.026	-0.01	0.248***	0.27***
	(0.019)	(0.029)	(0.022)	(0.030)	(0.044)
Regime 3	0.139	0.076	0.039	-0.221	-0.12
	(0.089)	(0.101)	(0.096)	(0.175)	(0.184)

Covariance & correlation matrix

4C

		QMJ	SMB	HML	UMD	MKT
Regime 1 *low-volatility*	QMJ	0.71 (0.1)	-0.42	-0.52	-0.5	-0.54
	SMB	-0.37 (0.04)	1.13 (0.18)	-0.07	0.43	0.45
	HML	-0.36 (0.03)	-0.06 (0.04)	0.7 (0.09)	0.29	0.15
	UMD	-0.47 (0.04)	0.51 (0.06)	0.27 (0.04)	1.26 (0.15)	0.52
	MKT	-0.71 (0.06)	0.75 (0.08)	0.2 (0.05)	0.92 (0.09)	2.51 (0.13)
Regime 2 *medium-volatility*	QMJ	0.89 (0.09)	-0.2	-0.04	0.54	-0.45
	SMB	-0.22 (0.05)	1.35 (0.1)	-0.06	-0.15	-0.04
	HML	-0.04 (0.04)	-0.07 (0.05)	0.97 (0.08)	-0.3	-0.3
	UMD	0.84 (0.08)	-0.29 (0.09)	-0.49 (0.07)	2.78 (0.18)	-0.36
	MKT	-0.85 (0.1)	-0.09 (0.1)	-0.58 (0.09)	-1.19 (0.17)	3.99 (0.26)
Regime 3 *high-volatility*	QMJ	5.08 (0.67)	-0.42	-0.11	0.4	-0.61
	SMB	-2.21 (0.46)	5.56 (0.72)	-0.15	0	0.25
	HML	-0.62 (0.32)	-0.84 (0.4)	5.92 (0.77)	-0.44	-0.03
	UMD	3.86 (0.66)	0.01 (0.73)	-4.62 (0.65)	18.26 (1.81)	-0.22
	MKT	-6.2 (0.97)	2.72 (0.69)	-0.36 (0.67)	-4.33 (1.47)	20.66 (2.04)
Regime 4 *low-volatility variation*	QMJ	0.28 (0.03)	-0.23	-0.29	0.15	-0.18
	SMB	-0.12 (0.02)	1.03 (0.06)	-0.11	-0.23	-0.42
	HML	-0.12 (0.02)	-0.09 (0.04)	0.6 (0.06)	-0.31	-0.47
	UMD	0.06 (0.02)	-0.18 (0.03)	-0.19 (0.03)	0.64 (0.07)	0.35
	MKT	-0.15 (0.04)	-0.68 (0.07)	-0.58 (0.06)	0.45 (0.05)	2.55 (0.14)

3C

		QMJ	SMB	HML	UMD	MKT
Regime 1 *low-volatility*	QMJ	0.44 (0.04)	-0.35	-0.36	-0.22	-0.34
	SMB	-0.24 (0.02)	1.04 (0.08)	-0.13	0.15	0
	HML	-0.19 (0.02)	-0.1 (0.03)	0.59 (0.05)	-0.03	-0.22
	UMD	-0.14 (0.02)	0.15 (0.03)	-0.02 (0.02)	0.86 (0.07)	0.39
	MKT	-0.34 (0.03)	0 (0.04)	-0.25 (0.04)	0.54 (0.05)	2.22 (0.09)
Regime 2 *medium-volatility*	QMJ	0.89 (0.09)	-0.19	-0.04	0.55	-0.43
	SMB	-0.21 (0.05)	1.38 (0.1)	-0.03	-0.18	-0.08
	HML	-0.04 (0.04)	-0.04 (0.05)	0.99 (0.08)	-0.3	-0.34
	UMD	0.87 (0.08)	-0.36 (0.1)	-0.51 (0.07)	2.84 (0.19)	-0.32
	MKT	-0.84 (0.1)	-0.2 (0.1)	-0.71 (0.1)	-1.13 (0.17)	4.3 (0.27)
Regime 3 *high-volatility*	QMJ	4.98 (0.66)	-0.42	-0.15	0.37	-0.62
	SMB	-2.16 (0.42)	5.37 (0.73)	-0.13	0.02	0.27
	HML	-0.79 (0.3)	-0.75 (0.37)	5.76 (0.76)	-0.41	0.01
	UMD	3.42 (0.57)	0.15 (0.65)	-4.12 (0.58)	17.36 (1.71)	-0.2
	MKT	-6.14 (0.89)	2.78 (0.64)	0.07 (0.61)	-3.81 (1.3)	19.94 (1.78)

2C

		QMJ	SMB	HML	UMD	MKT
Regime 1 *low-volatility*	QMJ	0.47 (0.03)	-0.29	-0.3	0.06	-0.33
	SMB	-0.2 (0.02)	1.04 (0.06)	-0.12	0.03	-0.04
	HML	-0.16 (0.02)	-0.1 (0.02)	0.62 (0.04)	-0.2	-0.25
	UMD	0.04 (0.02)	0.03 (0.03)	-0.16 (0.02)	1.07 (0.07)	0.14
	MKT	-0.35 (0.03)	-0.06 (0.04)	-0.31 (0.04)	0.23 (0.04)	2.43 (0.09)
Regime 3 *high-volatility*	QMJ	3.52 (0.3)	-0.38	-0.12	0.39	-0.6
	SMB	-1.42 (0.2)	3.93 (0.34)	-0.1	-0.02	0.21
	HML	-0.44 (0.14)	-0.4 (0.18)	4.01 (0.34)	-0.38	-0.06
	UMD	2.56 (0.28)	-0.12 (0.32)	-2.68 (0.26)	12.35 (0.81)	-0.22
	MKT	-4.24 (0.42)	1.55 (0.31)	-0.46 (0.29)	-2.93 (0.63)	14.31 (0.87)

Note: In the first block the transition probability matrices are shown, the second block displays the regime dependent means, while the last block displays the covariance and correlation matrices in each regime. The colored cells display the correlations. A darker red implies a larger positive correlation, while a darker blue implies a larger negative correlation. Standard errors are shown in parentheses. They are estimated with the OPG method. ***,**,* display significance at the 10%, 5%, and 1% level respectively, for a t-test testing the null that the regime-dependent mean return is 0. In the 2C model, moments have been placed under 'regime 3' to ease comparison.

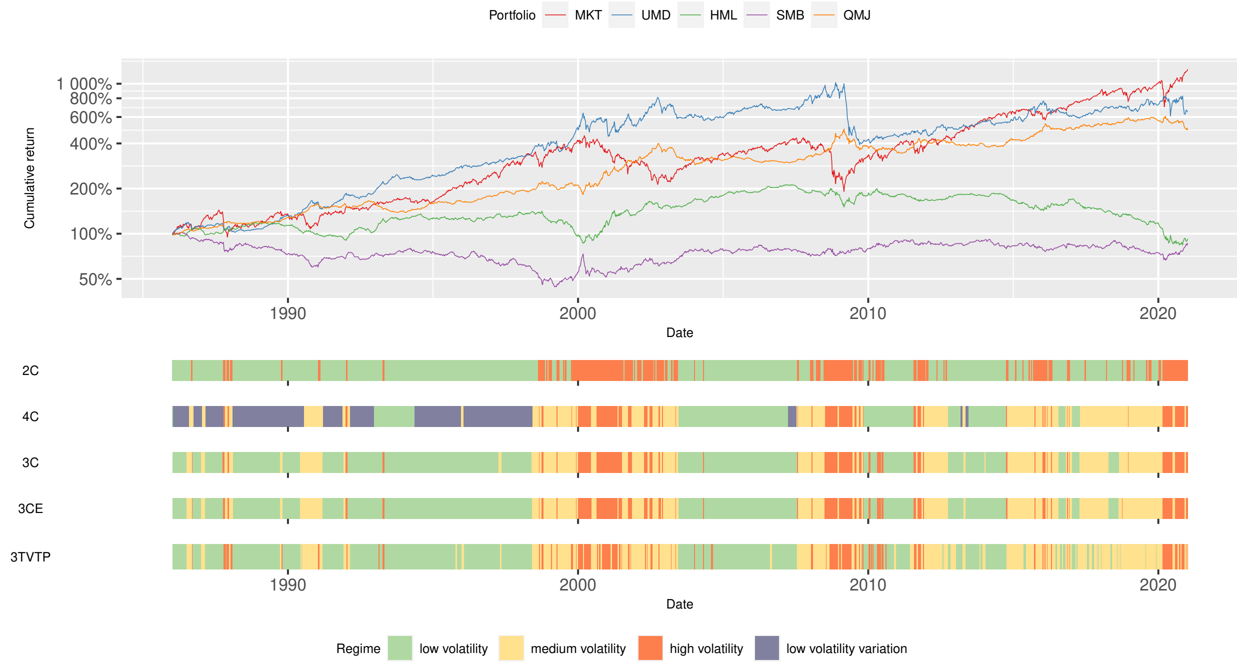

Figure 2: The cumulative returns of the factors, and sequences of most likely smoothed states of the models, estimated on the full sample. The top figure displays the cumulative returns of each of the factor portfolios, where the Y-axis is transformed to the log-10 scale. The bottom figure displays the sequence of regimes with the highest smoothed probability for the 2C, 3C, 4C, 3CE, and 3TVTP models.

of 2007.

The second, *medium-volatility*, regime exhibits lower persistence than the low-volatility regime, but with a transition probability to itself of 0.894, it is still reasonably high. Apart from transitioning to itself, transitions are most likely to the high-volatility regime. In terms of mean returns, the risk premia associated with UMD and QMJ are present, with significantly positive returns. Furthermore, HML has a significant negative risk premium, while the risk premium for SMB is not significantly present. Volatility in the second state is higher than in the first state, although considerably lower than in regime three. The 'ranking' of volatility of all the factors is the same as in all regimes: QMJ has the lowest volatility, while MKT has the highest.

The third regime – labelled as *high-volatility* – is much less persistent, with a 81.1% chance of remaining the same in the next period. Transitions to the medium-volatility regime are far more likely than transitions to the regime with low volatility; the chance of transitioning to regime three is 15.9%, in contrast to a chance of only 2.9% of a transition to regime one. The anomalies attributed to the QMJ and HML portfolios vanish in this state, with no mean return significantly different from 0. More so, the UMD portfolio attains a large significant effect in the opposite direction. Again, the risk premium of SMB is not significantly present. This is peculiar, because it implies that the size anomaly is not identified by the Markov-switching model in any of the states.

Although most returns cannot be significantly distinguished from 0, it is observed from the cumulative factor returns in Figure 2 that regime three coincides with many periods of downturns and large drops of MKT. In fact, many bear markets and market crashes of recent history fall in regime three, such as – among others – *Black Monday* in 1987, the burst of the *Dot-com bubble* in 2000, the *Financial Crisis* of 2007-08, and the recent *COVID-19-crash.* Therefore, when fitting this regime to the general known notions of market regimes, it can be considered as the classic 'bear' market.

Not only is there great variability in magnitude of volatility, the size and – even – sign of the risk premia between regimes, but also the correlations between factors differ. Besides differences in the strength of correlation occurring for all factors, the most variability is found for the MKT and UMD factors. In regime one, shocks to MKT are mildly negatively correlated with shocks to QMJ and HML, and have a slightly stronger positive correlation with UMD. This changes drastically in comparison to regime three, where the direction of correlation between UMD and MKT flips, correlation between HML and MKT vanishes, and SMB enters with a positive correlation. The picture is again very different in regime two, where all factors are negatively correlated with shocks to MKT. UMD also showcases switches in direction of correlation between regimes; shocks to UMD are negatively correlated to QMJ in regime one, while being fairly highly positively correlated to QMJ in the other regimes. The relation between UMD and SMB also varies greatly over the regimes; the relation is positive in regime one, negative in the second regime, and almost non-existent in regime three.

Although the specification with three states is chosen, it is not a clear-cut decision. Information criteria pointed at a four state model, and while the DST found less evidence against the null of correct specification for the three-state model, it was rejected nonetheless. Therefore, the models with a different number of states are discussed shortly below.

Comparison of the regime-dependent moments of the three models discloses an interesting feature: comparing the specification with two and three states, the regime-dependent moments differ greatly. This is in contrast with the moments of the three- and four-state model; the moments in state two and three remain mostly unaltered between the two models, while only estimates of regime one differ. In fact, it turns out that allowing for an extra state leads to a division of regime one into two more extreme regimes, that accommodate distinct returns and correlations for QMJ, SMB, and HML, while retaining approximately the same volatility. The restriction to three states leads to a state that can be seen as an 'average' of regime one and the – new – regime four, both in correlation and mean returns. From Figure 2 this idea of division of regime one is confirmed. The additional regime obtains the highest smoothed probability in periods that are fully spanned by regime one in the three-state model. These periods mostly occur at the very beginning of the sample. Periods where regime two or three have the highest smoothed probability of occurring are almost identical between the models.

Due to the similarity in two states for models 3C and 4C, and the dissimilarity in all states between 2C and 3C, I conclude that three states is the minimum number of states to model this process. Moreover, three states are actually preferred, because this model delivers regimes that occur throughout the whole sample, whereas the extra state of the four state model only appears important to the model fit in the beginning of the sample.

With the decision for three regimes, and a clear idea about the dynamics in this model with a low-, medium-, and high-volatility state, it remains to examine what effect explanatory variables have on this specification. Furthermore it remains to determine what the split of the 'average' returns and covariance matrix in Table 1 into three distinct cases brings about in portfolio optimization.

5.3 The Effect of Explanatory Variables

5.3.1 Implementation of Time-Varying Transition Probabilities

For estimation of the 3TVTP model, variable selection is performed with a step-wise procedure using the score test, as outlined in section 3.4.2. Implementation of this procedure however exposes an underlying numerical problem that becomes more severe when the number of states is increased: occasionally, a variable/ transition combination is selected due to it having an unreasonably inflated score statistic, paired with extremely slow convergence of the EM-algorithm. Due to the paired occurrence of these two problems, the cause of the problem is directly suspected to be related to the score function, which is important for both cases: in the score test, the conditional score is used for calculating the OPG estimate of the information matrix, and the regular score also appears in calculation of the statistic. In the EM-algorithm, the gradient of equation (15) is used by the BFGS algorithm to obtain a solution. This gradient turns out to be equivalent to the score of the log-likelihood in equation (9).

The extremely slow convergence most easily provides an answer to what the problem likely is: the likelihood surface for the transition is very flat. This suspicion is strengthened by the fact that problems only arise for variable/ transition combinations in transitions that rarely occur in the sample, and thus have very low transition probabilities anyway. Intuitively, this is explained by considering the S-shape of the softmax-function that is equation (2): in rarely occurring transitions, this function is very flat.

It also explains why the problems are especially prevalent when estimating a time-varying transition probability model with four states – or even more, but this has not been tested: inclusion of an extra state leads to more transitions that occur even less, due to states becoming more specific.

To circumvent the problem and to ensure that each selected variable significantly improves the model fit, an ex-post LR test is conducted after a variable has been selected. If the LR-statistic does not indicate a significant improvement on the 10% significance level, the variable/ transition combination is excluded and the step-wise procedure is continued with the variable with the next largest score statistic, if it surpasses the 10% significance threshold of course. Application leads to the sequence of selected variables and corresponding log-likelihood increases that can be found in appendix D.1.

For the three-state model, the just described problems did not occur frequently: only two times a variable/ transition combination was rejected. Besides the procedure above, I also standardize the explanatory variables by subtracting the mean and dividing by the standard deviation, to ease the numerical optimization procedure. This is only done for the TVTP model.

5.3.2 Implementation of the LASSO & regime-dependent return-predictors

To estimate the 3CE model using the LASSO implemented in the EM-algorithm, the regularization parameter λ first has to be chosen, as described in section 3.4.1. The λ sequence that is used here, together with model evaluation performance measures such as the DST and information critera, can be found in appendix D.2.

The disputed reputation of return prediction becomes apparent from the results in this table: no choice of λ leads to a lower BIC than the 3C model, nor does it lead to less evidence against the null of correct specification for the DST. Only the AIC provides any justification for addition of explanatory variables for return prediction: compared to the 3C model it is lower for the fully parameterized model with $\lambda = 0$ in each regime, as well as for the 6th and 7th set of λ's.

Based on the performance measures, the question arises if a model with regime-dependent return predictors should even be considered at all. I choose to do so, to illustrate the novel use of the LASSO within the EM-algorithm, and to further examine possible time-varying relations between predictors and returns. The 6th set of λ's is chosen, as this specification is accompanied by both the largest drop in AIC, and the largest increase in p-value for the DST, compared to a less restrictive set of λ's. This set, with λ equal to 0.101, 0.159, and 0.338 in regime one, two, and three respectively, also leads to a relatively parsimonious model, with four, six, and four predictors in regime one, two, and three respectively.

Table 6: The transition matrices and moments for the 3TVTP and 3CE model, estimated on the full sample.

Transition matrix

Three states with explanatory variables (3CE)

to \ *from*	Regime 1	Regime 2	Regime 3
Regime 1	0.977	0.029	0.024
Regime 2	0.012	0.896	0.156
Regime 3	0.010	0.081	0.815

Three states with time-varying transition probabilities (3TVTP)

to \ *from*	Regime 1	Regime 2	Regime 3
Regime 1	0.924	0.089	0.096
Regime 2	0.058	0.789	0.22
Regime 3	0.018	0.1155	0.687

Mean

3CE: See Table 13.

3TVTP:

	QMJ	SMB	HML	UMD	MKT
Regime 1	0.066^{***}	-0.013	0.051^{**}	0.196^{***}	0.274^{***}
	(0.022)	(0.035)	(0.026)	(0.031)	(0.05)
Regime 2	0.125^{***}	-0.008	-0.15^{***}	0.395^{***}	0.116
	(0.045)	(0.054)	(0.046)	(0.08)	(0.095)
Regime 3	0.125	0.069	0.156	-0.699^{**}	-0.104
	(0.146)	(0.168)	(0.16)	(0.281)	(0.307)

Covariance & correlation matrix

Three states with explanatory variables (3CE)

		QMJ	SMB	HML	UMD	MKT
Regime 1 *low-volatility*	QMJ	0.44 (0.04)	−0.35	−0.36	−0.22	−0.34
	SMB	−0.23 (0.02)	1.02 (0.08)	−0.13	0.16	0
	HML	−0.18 (0.02)	−0.1 (0.03)	0.58 (0.05)	−0.03	−0.22
	UMD	−0.14 (0.02)	0.15 (0.03)	−0.02 (0.02)	0.85 (0.08)	0.39
	MKT	−0.34 (0.04)	−0.01 (0.05)	−0.25 (0.04)	0.53 (0.05)	2.22 (0.11)
Regime 2 *medium-volatility*	QMJ	0.85 (0.1)	−0.17	−0.05	0.54	−0.43
	SMB	−0.18 (0.06)	1.29 (0.11)	−0.05	−0.16	−0.07
	HML	−0.04 (0.05)	−0.06 (0.06)	0.94 (0.09)	−0.31	−0.33
	UMD	0.82 (0.09)	−0.31 (0.11)	−0.49 (0.08)	2.73 (0.21)	−0.33
	MKT	−0.81 (0.11)	−0.17 (0.11)	−0.64 (0.1)	−1.11 (0.19)	4.11 (0.3)
Regime 3 *high-volatility*	QMJ	4.62 (0.81)	−0.39	−0.15	0.36	−0.62
	SMB	−1.93 (0.52)	5.15 (0.99)	−0.15	0.03	0.27
	HML	−0.76 (0.35)	−0.81 (0.48)	5.41 (0.96)	−0.42	0.03
	UMD	3.1 (0.67)	0.28 (0.86)	−3.96 (0.75)	16.38 (2.32)	−0.21
	MKT	−5.78 (1.06)	2.65 (0.85)	0.28 (0.77)	−3.67 (1.68)	18.85 (2.19)

Three states with time-varying transition probabilities (3TVTP)

		QMJ	SMB	HML	UMD	MKT
Regime 1 *low-volatility*	QMJ	0.41 (0.04)	−0.35	−0.38	−0.21	−0.31
	SMB	−0.22 (0.02)	1 (0.08)	−0.11	0.13	−0.07
	HML	−0.18 (0.02)	−0.08 (0.03)	0.56 (0.05)	−0.03	−0.24
	UMD	−0.12 (0.02)	0.11 (0.03)	−0.02 (0.02)	0.8 (0.07)	0.41
	MKT	−0.29 (0.03)	−0.1 (0.04)	−0.26 (0.04)	0.53 (0.05)	2.12 (0.1)
Regime 2 *medium-volatility*	QMJ	0.96 (0.1)	−0.19	−0.05	0.52	−0.45
	SMB	−0.22 (0.06)	1.31 (0.12)	−0.11	−0.14	0.03
	HML	−0.05 (0.04)	−0.13 (0.05)	1 (0.09)	−0.3	−0.28
	UMD	0.86 (0.08)	−0.26 (0.1)	−0.5 (0.08)	2.8 (0.2)	−0.33
	MKT	−0.9 (0.11)	0.07 (0.11)	−0.58 (0.09)	−1.12 (0.18)	4.17 (0.26)
Regime 3 *high-volatility*	QMJ	5.01 (0.68)	−0.42	−0.15	0.35	−0.61
	SMB	−2.22 (0.45)	5.69 (0.81)	−0.11	0.01	0.24
	HML	−0.81 (0.32)	−0.62 (0.4)	5.93 (0.82)	−0.41	−0.02
	UMD	3.31 (0.59)	0.07 (0.7)	−4.19 (0.61)	17.65 (1.84)	−0.18
	MKT	−6.26 (0.92)	2.59 (0.68)	−0.18 (0.67)	−3.53 (1.37)	20.79 (1.92)

Note: Estimates for the 3CE model are shown in the left column, estimates for the 3TVTP model in the right. In the first block the transition probability matrices are shown. The transition matrix for the model in the right column is not constant, therefore a reference matrix is shown, calculated as in equation (34). The transition probability coefficients can be found in Table 13. The second block shows the regime-dependent means. The last block displays the covariance and correlation matrices in each regime. The colored cells display the correlations. A darker red implies a larger positive correlation, while a darker blue implies a larger negative correlation. Standard errors are shown in parentheses. ***,**,* display significance at the 10%, 5%, and 1% level respectively, for a t-test testing the null that the regime-dependent mean return is 0.

5.3.3 Model Estimates

When explanatory variables are used to influence transition probabilities, the joint effect that they have on the transition probabilities is large. This can be seen by the reference transition probability matrix in the first block of the right column of Table 6. This matrix represents the weighted average transition probability over the whole sample, where each time-varying transition probability from state j to k is weighted by the filtered belief that the current state is j:

$$\bar{p}_{kj} = \frac{\sum_{t=1}^{T} p_{kj,t}\hat{s}_{t|t}^{j}}{\sum_{t=1}^{T} \hat{s}_{t|t}^{j}} \tag{34}$$

From the matrix it can be seen that transitions from one state to a different one occur with a higher probability than in the model with constant transition probabilities, on average: each state is less persistent than in the simple three state model, with larger off-diagonal transition probabilities as a consequence. On average, however, the same structure as in the simple Markov-switching model remains intact, where regime two functions as a transition state between regime one and three. Both regime one and three exhibit transition probabilities to regime two that are at least twice as large as the probability of transitioning to the other regime, on average. These effects also clearly show in the sequence of most likely smoothed states of the 3TVTP model in Figure 2: generally, the most likely smoothed states are similar to the 3C model, but the sequence is more fragmented, with regime switches occurring more frequently.

The size of the individual effect of the explanatory variables varies per transition, which can be seen by their marginal effects in appendix E.1. I choose to not discuss each of these effects separately, because one quickly becomes lost in their interpretation due to the multitude of transitions. Instead the economic significance of each of the variables is considered from the perspective of the mean-variance investor, by examining the ceteris paribus marginal effect of a standard deviation increase of a variable on the expected SR of the investor's portfolio. These effects will be discussed in one of the paragraphs below.

The second block in the right column shows that in comparison to the 3C model, also the regime-dependent mean returns change, but this difference is mainly limited to differences in the size of returns. There are no major differences in sign. The regime-dependent covariance matrices are extremely similar, as can be seen in the third block of the table.

The left column of Table 6 shows that when predictors of regime-dependent returns are allowed, the model is practically identical to the 3C model along all dimensions except the varying returns: the covariance matrices are very similar, and also the probabilities in the transition probability matrix are very close to the 3C model. Furthermore, the sequence of most likely smoothed states is also almost identical. Hence, the only dimension along which the model changes is in the endogenous regime-dependent returns, for which the coefficients of the variables can be found in appendix E.2. For the same reason as with the coefficients of the 3TVTP model, I also refrain from discussing these effects separately, and again focus on the marginal effect of a standard deviation increase on the expected SR.

The expected SR is considered because it offers a parsimonious interpretation of the effects that variables have on the investment opportunities of a mean-variance investor, both in the case when variables enter the model in the transition probabilities, and when they predict regime-dependent returns. Because the (expected) SR is a measure of risk-adjusted returns, it can be used as a gauge of the attractiveness of the investment climate of a mean-variance investor. An increase in the expected SR indicates a better investment climate.

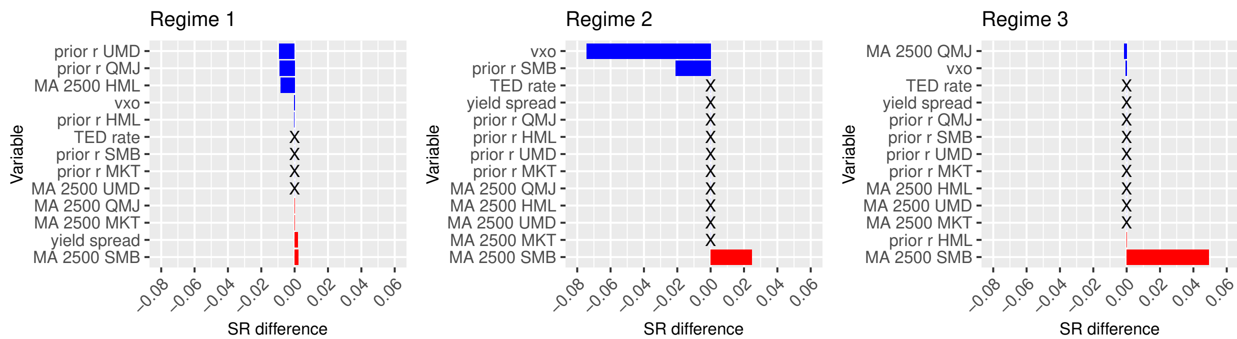

Figure 3: The effect in the 3TVTP model of a standard deviation increase of each of the variables, evaluated at their means, on the expected SR of a mean variance investor who faces the problem in equation (24), when the current regime is known. Effects are ceteris paribus. X's denote variables that are not selected for in either one of the transitions to the arrival states.

The marginal effect is considered when the current regime is known, and evaluated when each variable is at its mean, ceteris paribus. It is derived analytically, by calculating the moments using equations (22) and (23). In the 3TVTP model the moments vary by the predicted state beliefs $\hat{s}^k_{t+1|t}$, which change by the effects of each standard deviation increase in an explanatory variable that is considered. In the 3CE model, the moments vary due to changing $\hat{\boldsymbol{\mu}}_{k,t+1}$, by influence of the explanatory variables. Next, the portfolio optimization problem in equation (24) is solved in usual fashion, and the expected Sharpe ratio is calculated, now using estimated, instead of observed moments:

$$\hat{SR} = \frac{\hat{r}^{p,(e)}}{\sqrt{\boldsymbol{w}'\hat{\boldsymbol{\Sigma}}\boldsymbol{w}}}. \tag{35}$$

Figure 3 shows that in the 3TVTP model, effects on the expected SR evaluated at the mean value of all variables are relatively small in the low-volatility regime one, in comparison to the other regimes. In these other regimes it stands out that only very few variables considerably affect the expected SR. *MA 2500 SMB*, is one of these variables, with a positive impact on the expected SR in the medium-volatility regime, and a larger positive effect in the high-volatility regime. This variable also has a small positive impact in regime one. It appears that its effect grows with the volatility of the regime, and that a high valuation of the size factor indicates good expected investment conditions, from the perspective of a mean-variance investor. Another derivative of *SMB* returns, namely its prior one month return has a negative effect on the expected SR. Lastly, *vxo* has a large negative impact on the expected SR in the medium-volatility regime two.

For the variables in the 3CE model, Figure 4 shows that *vxo* affects the expected SR in a positive manner in the low-volatility regime, while this relation is inversed in the high-volatility regime three. Hence, an increase in implied volatility can be interpreted as a positive signal for investment

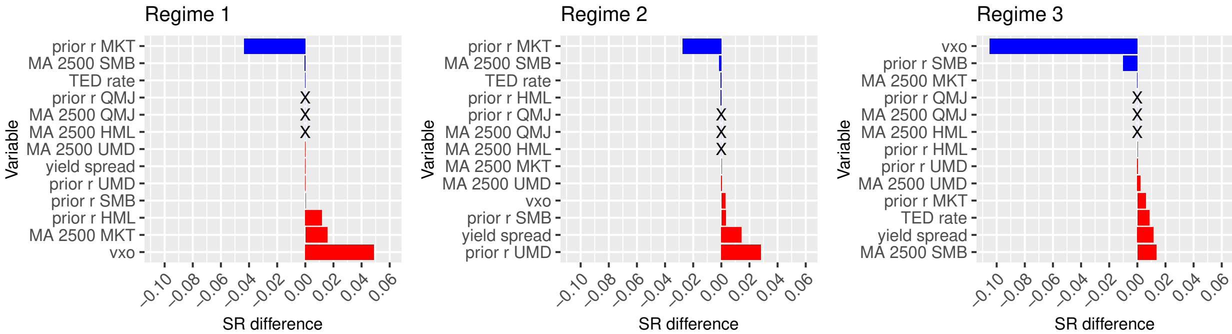

Figure 4: The effect in the 3CE model of a standard deviation increase of each of the variables, evaluated at their means, on the expected SR of a mean variance investor who faces the problem in equation (24), when the current regime is known. Effects are ceteris paribus. X's denote variables that are not selected by the LASSO.

climate when volatility is low, but as a negative signal when volatility is high. Furthermore, *MKT* momentum negatively impacts the expected SR in both the low- and medium-volatility regime, while its effect is marginally positive in the high-volatility regime. Momentum of *UMD* – therefore 'momentum of momentum'– positively affects the expected SR in the medium-volatility regime. Also the yield spread has a reasonably large positive effect on the expected SR in both the medium- and high-volatility state, but this effect is not present in the low-volatility regime.

The differences in effects between the two models are remarkable. For one, the effects show that *vxo*, although being the only variable that has a large impact in both models, has a completely different impact on the expected SR through its influence on the regime than through the returns in the regime. Furthermore, momentum, of the market specifically, is one of the main drivers of differences in expected SR when considered as an predictor of regime-dependent returns, but nowhere to be found as a determinant for the probability of transition. Conversely, *MA 2500 SMB* exhibits much stronger effects through its impact on the transition probabilities, than by its appearance as a return predictor.

5.3.4 Model Performance

When explanatory variables are allowed in the specification with three states, there is a large discrepancy in the values of performance measures, dependent on which part of the model the explanatory variables enter.

The rightmost column in Table 2 shows that allowing transition probabilities to vary increases model fit. Both information criteria are lower in comparison to the 3C model, and the p-value of the multivariate DST is a factor 10 larger, although there still is just enough evidence to reject the null at the 1% level. Be that as it may, Figure 1 shows that the 3TVTP model is the only model

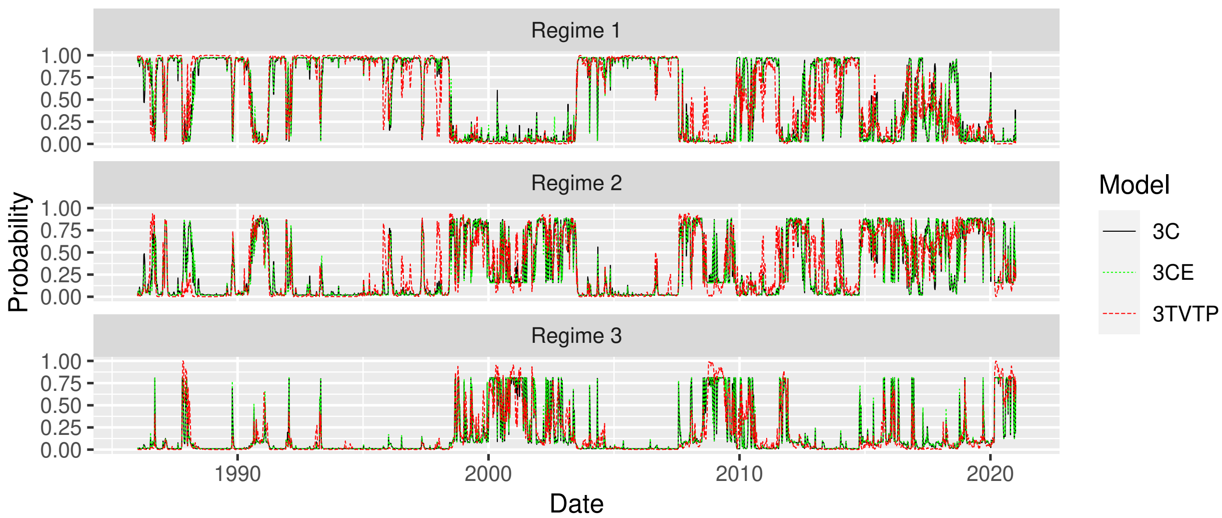

Figure 5: The predicted state probabilities estimated on the full sample for the 3C, 3TVTP, and 3CE model.

of which the z^* actually approach the desired standard normal distribution, in terms of skewness and kurtosis. Judging from the RMSE in Table 4, the better fit of the model and density does not improve point return prediction at all.

The RCM is higher than for the 3C model, where in general a low RCM is desired. However, it should be noted that a low RCM does not necessarily imply that the switches are correctly predicted, but only that they are sharp, with a small frequency of periods of uncertainty on the nature of regimes (Guidolin, 2011). By allowing the transition probabilities to be endogenous, it is expected that there is much more variability in them, with a higher RCM as a result.

By examining the p-values of the univariate DST in Figure 3, it becomes clear that the improvement in model fit is largely thanks to improvements in the specification of the density of QMJ; there is not enough evidence to reject the null of a correct specification of the density for this factor. In addition, there is less evidence against the null for SMB and UMD, while there is more evidence found against the null for MKT and HML. I however still fail to reject the null of correct specification of the density of HML at the 5% significance level.

Because I established that the identified regimes are practically identical to the 3C model in terms of correlation and volatility, the improvement in fit, and decrease in evidence against a correctly specified density, has to be caused by either the different estimates of regime-dependent mean returns, or by more accurate predicted states, which will be discussed below. It is hard to determine which one of these two differences causes the improvement in fit, but either way the time-varying transition probabilities are involved.

Figure 5 shows that the 3TVTP model differentiates itself from the 3C model by being more confident, barring its generally higher fluctuation, in its predicted regime beliefs: in periods where all models have a high (low) predicted state belief, the predicted beliefs from the 3TVTP model nudge just a bit more towards the limit of 1 (0). It turns out that this difference in confidence disappears when beliefs are updated to their smoothed versions. Then, beliefs of all models are often on the limits, as can be seen in appendix F.1. The distinction between predicted and smoothed beliefs is made here, since predicted beliefs are used in the DST, and therefore are relevant in explaining the improvement in density fit.

As was already briefly discussed in the section on implementation of the LASSO, there is not much evidence for an improvement in model fit when regime-dependent return predictors are included. Both BIC and and DST do not favor the 3CE model over the 3C model, only AIC is slightly better while the RCM obtains a similar score.

Similar to the 3TVTP model, correlation and volatility in each regime of the 3CE model are almost identical to the 3C model. Additionally, the same holds for the predicted beliefs, by Figure 5. This resemblance in predicted beliefs between the 3C and 3CE model reveals that the states are mostly identified by their volatility; differing returns have little impact on the predicted state beliefs. Furthermore, the same logic applied to the reasoning in the 3TVTP model now flows the other way: by similarity of correlation, covariance, and predicted state beliefs in combination with worse model fit, I conclude that there is ample evidence that regime-dependent predictors do not appear in the correct density of joint factor returns.

With the conclusion in mind it is striking that the 3CE model significantly adds economic value in-sample: it attains a significantly higher Sharpe ratio than the 3C model, or any other model for that matter. It achieves this thanks to a considerably higher realized return. To find a cause for this difference, a deeper dive in the return estimates of all models is needed. After all, a difference in return prediction is the only factor that can cause the dichotomy in portfolio performance between the 3C and 3CE model, as I have just shown that regime beliefs and regime-dependent covariances are practically identical.

While the RMSE of the 3CE model do appear to be roughly 1% lower than those of the other models, it is hard to believe that this small improvement in point prediction is responsible for the large increase in portfolio returns. Therefore another cause is sought. In appendix F.2 it can be seen from the cumulative returns of applying the models in an in-sample portfolio optimization context that the 3CE model does not achieve its higher returns by a particular period of outperformance. Instead, its returns appear to consistently be a factor larger than those of other models. Therefore I deduct that the 3CE model systematically does something better than the other models.

The most probable explanation found is that the 3CE model is able to jointly predict the direc-

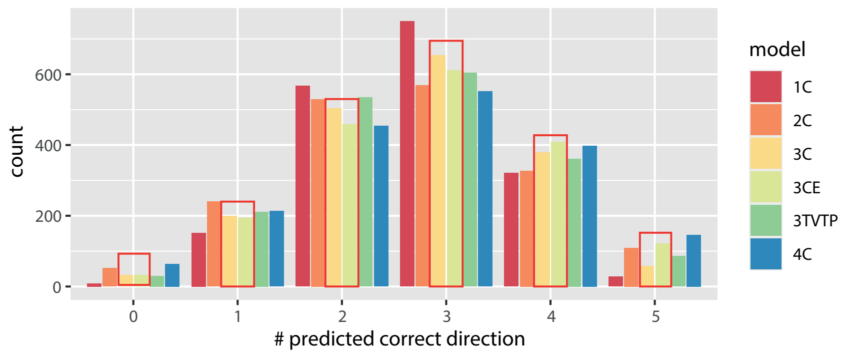

Figure 6: The number of times each models correctly predicted the direction of factor returns. The counts for the 3C and 3CE model are outlined in red.

tion of returns better than the 3C model. Even though it is a massive simplification to imply that superior sign prediction instead of point prediction provides an explanation, it is deemed likely, due to the structure of the (mean-variance) problem with no borrowing- and short-selling restrictions. Due to the restrictions, each weight is even more dependent on the other weights than in normal mean-variance optimization. A bad prediction therefore does not only affect the weight of a factor itself, but also is detrimental to other weights. This is why the 'joint' part of prediction is thought important. The sign is deemed important because when the predicted direction is correct, the portfolio still 'utilizes' the factor and its relations to other factors in the right manner, but just in a different magnitude. Table 15 in appendix F.3 shows that there is indeed a positive correlation between the number of correctly predicted directions and portfolio returns

Figure 6 illustrates the just made argument: comparing the 3C and 3CE model – which are outlined by a red rectangle – it can be seen that the 3CE model correctly predicts the direction of all five factors more than twice as often as the 3C model. It further correctly predicts four out of five factors slightly more often, while also predicting only two out of five correctly less frequently.

Comparison with other models by this figure is less straightforward, since they rely on different number of states and covariance matrix estimates (2C, 4C), or different predicted state beliefs (3TVTP). Still, the mean of the number of correctly predicted factor returns provides a performance measure that monotonically increases with portfolio return: this measure equals 2.71, 2.66, 2.72, 2.84, 2.72, and 2.79 for the 1C, 2C, 3C, 3CE, 3TVTP, and 4C model respectively.

To summarize the discussion on explanatory variables, a model with regime-dependent return predictors seems to economically benefit from the joint direction that the predictors provide, but a model with constant regime-dependent mean returns and high noise is more likely to model the density. Furthermore, I find some evidence that explanatory variables drive this process through influence on the transition probabilities. The framework of three regimes identified by their increasing

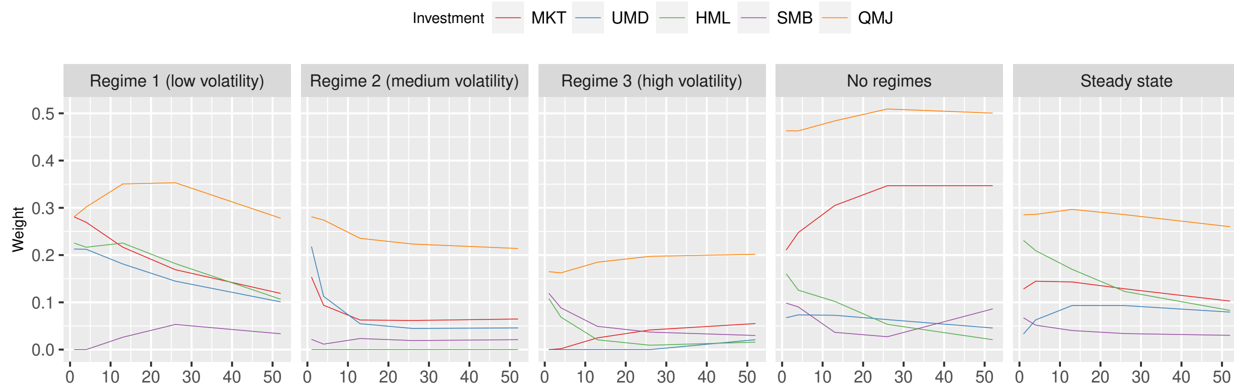

Figure 7: The weights of the optimal solution of the portfolio optimization problem in equation (24) for multiple investment horizons, when the current regime is known, together with the weights of the 1C model (No regimes), and those when state beliefs are at their long-term equilibrium (Steady state). $1 - w_{MKT}$ is invested in the risk-free asset. Moments of the Markov-switching models are calculated using algorithm 1 with regime-dependent return vectors and state covariance matrices estimated with the full-sample estimates from the 3C model.

volatility remains intact, after including explanatory variables in any way.

5.4 The Effect of Regimes on Investment Decisions

Because the three-regime model identifies regimes with completely different returns and covariance matrices, the prevailing regime has a large effect on the weights that a mean-variance investor allocates to the factors. In addition, because the model incorporates uncertainty about the future states, these weights vary greatly when longer investment horizons are considered, giving rise to intertemporal hedging demands. These findings are illustrated in Figure 7, where the weights for each factor are displayed for a buy-and-hold mean-variance investor for multiple investment horizons, when the current regime is known, and estimates from the 3C model are used.

When the current state is the *low-volatility* regime one, each factor receives a large weight, with the exception of SMB. As mean returns are positive, and volatility low, this is an attractive investment climate for a mean-variance investor. Consequently, (s)he invests the full allocation in the risky portfolios – the discussion of most of them being a zero-cost portfolio aside – when a short investment horizon is considered. Because regime one is highly persistent, the investor can be confident that a transition to a less favorable regime is not very likely in the near future. Weights therefore adapt slowly to other levels when the investment horizon increases, preparing for worse times yet to come.

The *medium-volatility* regime two is less persistent, and the probability of transitioning to the generally unfavorable regime three is much larger than from regime one. This leads to a rapid decrease in optimal weights between the single- and four-week investment horizon for MKT and UMD. These factors attain a high positive mean return in regime two, but a large negative mean return in regime three, pressing expected multi-period returns downwards. In this regime the investor does not fully invest into the portfolios anymore: with a single-period investment horizon about two-thirds are invested in the portfolios. This amount shrinks to roughly one-third with a year-long horizon.

In *high-volatility* regime three, the total amount invested in the portfolios hovers at around one-third for all investment horizons, but the composition of the allocation does vary with horizon time. For shorter horizons HML and SMB make up a a large part of the risky part of the portfolio, while none of it is invested in UMD and MKT. The weight on the former decreases, and the weight on the latter however increases with a longer investment horizon. Now the model recognizes that less volatile states are likely in the future, in which UMD and MKT thrive, while SMB slumps.

Comparing the weights of the model without regimes to those of the 3C model when the current regime beliefs are the long-term equillibrium probabilities of each regime, it can be seen that interestingly, allocations are not nearly the same. Using estimates from the Markov-switching model leads to weights that are more precautious. In addition, using estimates from a Markov-switching model leads to less risk-taking as the horizon increases.

The effects that regimes have on investment decisions is now clear, when the current regime is known. This last distinction is of importance, because in practice, this is not the case. Instead, the model attaches a probability to the prevailing state. Another barrier for implementation in practice is that the model has to provide estimates for unknown data. So far, all relations deducted are on data that the model was also trained on. To fully answer the remaining question what effect acknowledgement of regimes has on the investment decisions and performance of a mean-variance investor, this practical case has to be considered. This is what is done in the next section on OOS performance.

5.5 Out-of-Sample Portfolio Performance

When comparing the one period ahead OOS performance of the three-regime model, its variations, and the model that is ignorant of regimes in Table 7, a first observation is that all Sharpe ratios are unsurprisingly lower than in-sample. What does come as a surprise, is that there is a major difference in the hierarchy of the models in terms of performance, in comparison to their in-sample counterparts. Moreover, mutual differences in performance between Markov-switching models are substantially larger than they were in-sample.

Most notably, superior performance of the 3CE model vanishes completely. It is now the worst

Table 7: OOS performance measures of the optimal portfolio solution from equation (24) for one period-ahead returns, using inputs of the 3C, 3TVTP, 3CE, and 1C model, estimated with the expanding window approach.

Model	*Performance measure*				*Sharpe ratio difference*		
	$\bar{r}^e$ (%)	σ (%)	Utility (*10,000)	SR (annualized)	1C	3TVTP	3CE
3C	0.075	0.506	9.527	1.072	0.335	-0.12	0.404
					(0.147)	(0.508)	(0.244)
3CE	0.076	0.817	8.915	0.669	-0.069	-0.523	
					(0.829)	(0.147)	
3TVTP	0.091	0.552	11.043	1.192	0.455*		
					(0.051)		
1C	0.054	0.53	7.387	0.739			

Note: Moments of the Markov-switching models for input in equation (24) are calculated using predicted state beliefs ($\hat{\boldsymbol{S}}_{t+1|t}$) as in equations (22, 23). The return vector and covariance matrix of the model which does not take regimes into account (1C) are the historical mean and covariance up to each of the 10 splits of the expanding window. The weekly mean excess portfolio return, standard deviation, realized utility, and Sharpe ratios are reported. Weekly Sharpe ratios have been annualized by multiplying by $\sqrt{52}$. *P*-values for the test of differences in Sharpe ratios are in parentheses. They are calculated using the test from Ledoit and Wolf (2008). This test is conducted on the non-annualized Sharpe ratios. ***,**,* display significance at the 10%, 5%, and 1% level respectively.

performing model in terms of risk-adjusted returns. This is mainly due to its large variance. Figure 8 shows that this is caused by greatly varying – mispredicted – weights. It is clear that the ability of the 3CE model to superiorly predict the joint direction of returns does not hold OOS.

The next worst performing model is the model that does not account for regimes, and therefore uses the mean and covariance matrix of historical returns up to the latest date in the training sample as inputs for portfolio optimization. A portfolio formed by this model suffers from lower average portfolio returns. The root of this problem lies in its inability to adjust its weights accordingly in periods of economic downturn. By the cumulative returns of all models in the OOS period in Figure 9 it can be seen that the 1C model suffers large losses during the crisis of 2007, in contrast to the Markov-switching models – with the exception of the jumpy 3CE model – which experience practically no fallback during this period. From the weights of the 3C and 3TVTP model in Figure 8 it can be seen that portfolios formed by the models achieve this favorable characteristic by assigning a considerably smaller part of the allocation to the risky factors and more to the riskless asset, in periods that the models classify as the high-volatility regime. Taking on less risk in these periods often is the smart move, because they mostly interfere with periods of economic downfall. There is however a drawback: because of their risk-averse behaviour in periods of high volatility, the models benefit not nearly as much from the sharp recovery after the Covid-19 crash as the 1C or 3CE model. Still, both the 3C and 3TVTP model achieve higher cumulative returns, paired

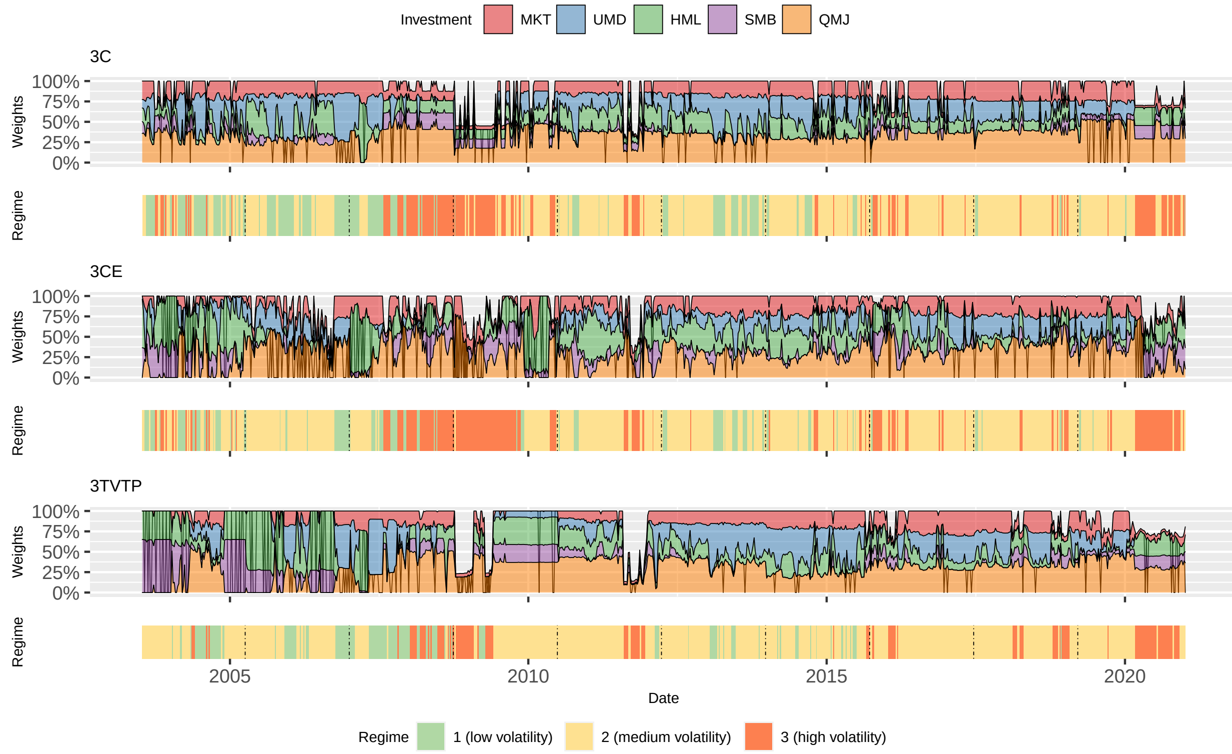

Figure 8: The OOS portfolio weights and the sequences of most likely smoothed states of the 3C, 3CE, and 3TVTP model, estimated with the expanding window approach. Each of the top figures displays the OOS weights from the optimal solution of the portfolio optimization problem in equation (24) for one-period ahead returns. $1 - w_{MKT}$ is invested in the risk-free asset. Moments for input in equation (24 are calculated using predicted state beliefs ($\hat{\boldsymbol{S}}_{t+1|t}$) as in equations (22, 23). Each of the bottom figures displays the sequence of most likely smoothed states. The dotted lines are placed at the start of each new division in the expanding window, and thus where the estimates are recalibrated.

with higher Sharpe ratios, than a model that does not acknowledge regimes.

The difference in Sharpe ratio with the 1C model however only is significant for the 3TVTP model, by the Ledoit and Wolf (2008) test. It is driven by a larger mean portfolio return than the other models. For the source of this higher return I do not have to dig as deep as in the in-sample case. The cumulative returns in Figure 9 namely show a period of abnormally high returns by the 3TVTP model starting in August 2007 and ending in December 2008.

The sequence of most likely smoothed states in Figure 8 reveals that the 3TVTP model differentiates itself from the the 3C and 3CE models by more frequently classifying regime one instead of three as most likely during this period. Inspection of the OOS predicted states – which are used in finding the optimal weights – of each model during this period in Figure 10 shows that the model further attaches a lower probability to the high-volatility regime, in cases where it is most likely. Since regime one is a low-volatility state, and regime three a high-volatility state, the

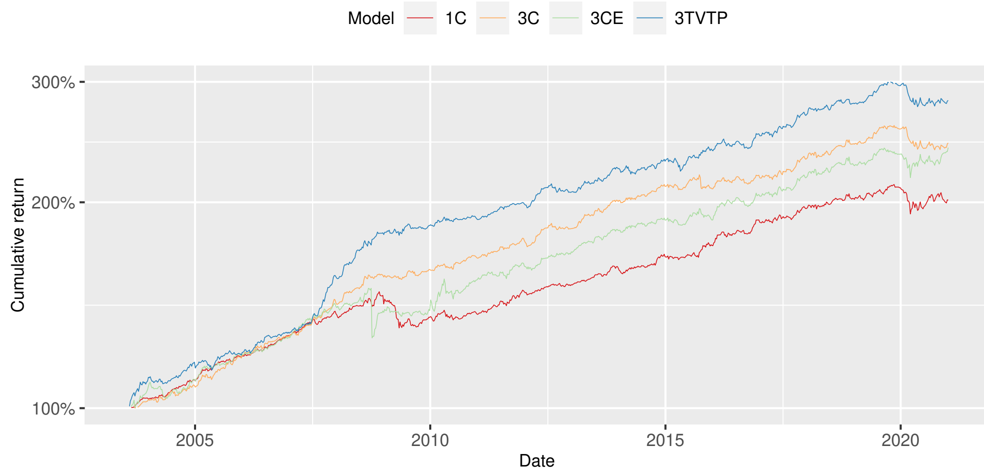

Figure 9: The OOS cumulative returns from the portfolios formed by solving the portfolio optimization problem in equation (24) for 1-period ahead returns using inputs of the 3C, 3TVTP, 3CE, and 1C models, estimated with the expanding window approach. Moments of the Markov-switching models for input in equation (24) are calculated using predicted state beliefs ($\hat{\boldsymbol{S}}_{t+1|t}$) as in equations (22, 23). The model which does not take regimes into account (1C) takes the historical mean and covariance of h-period returns up to each of the 10 splits of the expanding window as input for the portfolio optimization problem. The Y-axis is transformed to the log-10 scale.

3TVTP model takes more risk as a consequence.

Subsequently, a portfolio formed with estimates of this model fully allocates its wealth in the risky assets, whereas a portfolio from the 3C model does not, judging from Figure 8. Furthermore the distribution of the weights is different, with periods where a large weight is invested in UMD, and a larger weight on QMJ in general. Rightfully so, as these factors performed well in this period, as Figure 2 shows. By the superior performance thanks to distinct regime-estimates, I therefore confirm the hypothesis that time-varying probabilities lead to better state prediction OOS.

Figure 10: The predicted beliefs of each of the three-state models in the period of outperformance by the 3TVTP model.

Table 8: OOS performance measures of the optimal portfolio solution from equation (24) for h-period ahead returns, using inputs from the 3C and 1C model, estimated with the expanding window approach.

Model	*Performance measure*				
	$\bar{r}^e$ (%)	σ (%)	Utility (*10,000)	SR (annualized)	Sharpe ratio difference
			4 week horizon		
3C	0.210	0.666	29.926	1.138	0.726*** (0.000)
1C	0.178	1.563	10.168	0.412	
			13 week horizon		
3C	0.559	1.092	85.552	1.024	0.553*** (0.000)
1C	0.700	2.973	-476.362	0.471	
			26 week horizon		
3C	1.012	1.714	160.269	0.835	0.378*** (0.002)
1C	1.380	4.269	-5008.031	0.457	
			52 week horizon		
3C	1.833	2.653	300.782	0.691	0.237** (0.048)
1C	2.663	5.872	-57452.001	0.454	

Note: Moments of the 3C model for input in equation (24) are calculated using algorithm 1. The return vector and covariance matrix of the model which does not take regimes into account (1C) are the historical mean and covariance of h-period returns up to each of the 10 splits of the expanding window. The h-period mean excess portfolio return, standard deviation, realized utility, and Sharpe ratios are reported. Weekly Sharpe ratios have been annualized by multiplying the 4-, 13- and 26-period returns by $\sqrt{13}, \sqrt{4}, \sqrt{2}$ respectively. P-values for the test of differences in Sharpe ratios are in parentheses. They are calculated using the test from Ledoit and Wolf (2008). This test is conducted on the non-annualized Sharpe ratios. $^{***},^{**},^{*}$ display significance at the 10%, 5%, and 1% level respectively.

While for a one-period investment horizon only the 3TVTP model significantly outperforms the 1C model in terms of risk-adjusted returns, the 3C model also achieves this feature when longer horizons are considered. As can be seen in Table 8, the 3C model achieves a significantly higher Sharpe ratio than the model without regimes for all multi-period investment horizons. More complex models using explanatory variables are not considered, because these variables would then also have to be predicted.

Its performance is strongest for horizons that lie in the foreseeable future, e.g. four weeks or one quarter of a year. Whereas superior Sharpe ratios where mostly driven by higher returns for the one-period investment problem, for longer periods they are attained by keeping volatility low: for the 4-, and 13-week horizons it is almost three times as low as in the 1C model. This very clearly showcases the strength of the Markov-switching model to take future periods of high volatility into account, as was discussed in the previous section. The allocations delivered by the 3C model hedge away some of the risks of these future periods of economic downturn.

The figure in appendix F.4 shows that the model is able to deploy the same mechanism OOS as was found using the full sample: as the investment horizon increases, the amount of weight that is invested in risky assets decreases.

6 Robustness Tests

To examine the stability of the results, multiple checks of robustness have been conducted.

First of all, each factor is modeled as a univariate Markov-switching model with three states, to determine if it is not too restrictive to force the factors to have three regimes that all reside in the same periods. It may well be that the regimes of each factor are uncorrelated to those of others. Figure 14 in appendix G.1 shows that this however is not the case for most factors: generally the most likely smoothed states of each single factor abide by the same periods. There are some exceptions for this rule, where regime three is identified instead of regime one. This is prevalent for the latter part of the sample period for QMJ, but the most severe mismatch is with MKT.

The moments of the single factor Markov-switching models in Table 16 further show that when estimated separately, regimes are more distinctive, with more extreme moments. This especially shows in the returns. The estimates on regime-dependent volatility are similar to the multivariate model – albeit more extreme – with a ranking from low-volatility regime one to high-volatility regime three.

Secondly, the portfolio optimization problem is solved with the no-short-selling constraint substituted by the following constraint:

$$w_{i,t} \geq -1 \qquad \text{for } w_{i,t} \in \boldsymbol{w}_t. \tag{36}$$

This is done due to the construction of the investment factors. Since each factor is a zero-cost portfolio that is long one leg and short another, it may seem odd to subsequently restrict an investor to short-sell this portfolio while this is essentially just flipping the short and long leg. Of course, these factors are not products that one could actually invest in, but it is good to test whether a slightly different problem causes major differences in results.

All OOS results are reproduced for this altered problem in appendix G.2. From Figure 15 and Table 17 it can be seen that the result of the 3TVTP model performing best is robust to the different problem formulation for one-period returns, both in terms of cumulative returns and risk-adjusted returns. By way of contrast, the 3C model performs much worse for a one-period investment horizon. Nonetheless, it retains its significantly higher risk adjusted return in comparison to the 1C model for any of the multi-period investment horizons, as is shown in Table 18.

Thirdly, the consistency of the OOS estimates is checked, of which the results can be found in appendix G.3. In Figure 16, where the OOS estimates of moments based on each of the 10 vintages are displayed, it can be seen that estimates of mean returns are stable for both the 3C and 3TVTP model in regime one, and that instability increases – logically – with the volatility attributed to a regime. As a result, estimates of the mean in regimes two and three are not very robust to a changing training period; estimates experience sign changes and the ordering of mean returns – which remains preserved in the estimates of mean returns in regime one – is mixed up.

To compare the evolution of the OOS covariance matrices, a simplification is made to only look at the volatility of each factor, for conciseness. Estimates for the volatility are much more robust, regardless of the regime. The internal ordering remains intact for each training period, and the only large altercation is found in the volatilities of regime three, of which the level increases after the credit crisis is included in the training set.

Finally, the consistency of the variable selection procedures is inspected. Heatmaps of the amount of times a variable is selected by the iterative score test for the 3TVTP model, and the LASSO for the 3CE model, can be found in Figures 17 and 18 respectively in appendix G.4. In a perfectly robust variable selection procedure, one would have 10 appearances in the models on the positions that have coefficients in Tables 13 and 14. Clearly this is not the case. This may be because the *MA 2500*- and *prior r*-variables are relatively highly correlated to other variables of their kind, as can be seen in the correlation matrix of Table 10 in appendix B. Therefore a variable selection procedure may select one variable for a certain regime or transition for one training set, while choosing a different one in the next.

There are only a few consistencies with the full sample results of variable selection. That is that *vxo* is considered important in both models, with a large number of appearances. The same holds for *prior r MKT* in the 3CE model.

7 Conclusion & Discussion

In this thesis, I examine what specification of a Markov-switching model describes the dynamics of factor returns best, and what these dynamics are. Furthermore, I inspect how incorporating multiple regimes in a portfolio optimization problem affects the investment decisions and performance of a mean-variance investor who manages a portfolio of factors, in comparison to an investor who is ignorant of regimes. Lastly, I investigate whether letting explanatory variables predict either transition probabilities, or the regime-dependent returns, enhances the model fit and investment performance.

My conclusion is that a Markov switching model with constant regime-dependent returns, three states and time-varying transition probabilities models the joint dynamics of factors most properly, because it offers the best combination between parsimony, model fit, and economic interpretation.

I find that this model, but also Markov-switching models in general, are able to enhance investment performance of a mean-variance investor, mainly thanks to the ability of Markov-switching models to robustly identify regimes with different volatilities, both in-sample and OOS. As a consequence, the model is able to invest more cautiously in periods of high volatility. The role of regime-dependent mean returns seems limited, because the signal-to-noise ratio is very high and point predictions of returns are therefore poor. In line with this conclusion, I also find that predictors for regime-dependent returns do not enhance the model; some predictability is found in-sample, but it vanishes OOS. On the other hand, predictors for time-varying transition probabilities are a valuable inclusion to the model in the context of asset allocation. A model with time-varying transition probabilities provides better estimates of the predicted regime-probabilities OOS, before information about current or future observations is considered, with better assessment of the prevailing (co)variance structure as a result. I further find that superiority of Markov-switching models is even more prominent when multi-period investment horizons are considered. A simple three-state Markov-switching model significantly outperforms a model without regimes on multiple investment horizons, because it can take the possibility of future periods of adversity into account. Consequently, a portfolio formed by using this model achieves considerably lower volatility, with higher risk-adjusted returns and utility as a result.

The result of three regimes being the optimum for the joint prediction of factor returns is contradictory to the findings of Guidolin and Timmermann (2008), which is the only paper that is very closely related to the subject examined here. They find four regimes to be optimal for the joint process of the size, value and market factors. The discrepancy results from differences in model selection criteria. They base their choice solely on information criteria, and accept a fourth regime that appears only in very select periods in the beginning of their sample. In this thesis, the choice is made to attribute more importance to parsimony and economic importance. Furthermore the DST is seen as the most important measure of model fit. If only AIC and BIC would have been used, the specification with four states would similarly have been found optimal.

This is however not the only result that I find that contradicts their findings: the identified regimes are very different in terms of mean returns, and correlations. This may be because their sample period starts in the 1930s, and does not include the two major crises that occurred after 2007. Nevertheless, this reconciles with the finding in my thesis that mean regime-dependent returns are not robust on the training set.

On the other hand, there is a similarity in findings, namely the identification of a high-volatility state and two similar low-volatility states, which was found here for the four-state model. It turns out that identification of such a high-volatility state is not limited to Markov-switching models for factor returns. Regardless of the asset class(es) under investigation, many papers find regimes that are identified based on their volatility (Ang and Bekaert, 2004; Guidolin and Timmermann, 2008; Kole and Van Dijk, 2017). In contrast to the finding of Ang and Chen (2002) for U.S stock returns, this thesis finds that correlations between factors and the market are not much higher in this high-volatility regime.

The result that regimes are mostly determined by their covariance matrix is in line with findings of Kole and Van Dijk (2017), who explicitly test this matter by estimating a Markov-switching model with variance that is restricted to be constant between regimes. They find that this model is not able to properly distinguish bull and bear markets of stock returns.

The role of volatility in an asset allocation context is also found similar in other research: Ang and Bekaert (2004) find that their Markov-switching model achieves superior performance by its ability to switch to cash in periods of high volatility.

Some of the limitations in this thesis are related to the utilization of the results in reality. This begins with the data that is used: the factor portfolios are not actual investable assets. Instead, ETFs could be used, but not enough data was found. In addition, the effects of transaction costs should be taken into account. Lastly, it may well be that the potential of the Markov-switching model is undersold because the expanding window only uses 10 vintages due to computational constraints, instead of updating at each time t. Further research on actual implementation of Markov-switching models for factor timing could follow up on either one of these points, but the last point is most interesting, as it would require a more efficient algorithm.

Another point of further research, that I personally would find most interesting, would be to examine the properties of the LASSO in the EM algorithm. For one, ideally the LASSO would have been implemented not only for the regime-dependent return predictors, but also for the variables appearing in the transition probabilities. However, no easily applied solution was found within the scope of this research. Another possible path of further research would be to examine if the optimization of the regularization parameter could be integrated more elegantly into the EM algorithm. The solution implemented in this paper based on 10 sets of three regularization parameters was very limited, and it is likely that a better performing model could have been found, had there been a way to optimize the regularization parameter in each regime separately, instead of jointly.

Lastly, it would be interesting to examine how the Markov-switching model holds up against more complex baseline models, both in terms of density specification and portfolio performance.

References

Ammann, M. and Verhofen, M. (2006). The effect of market regimes on style allocation. *Financial Markets and Portfolio Management*, 20(3):309–337.

Ang, A. (2014). *Asset management: A systematic approach to factor investing*. Oxford University Press.

Ang, A. and Bekaert, G. (2002a). International asset allocation with regime shifts. *The Review of Financial Studies*, 15(4):1137–1187.

Ang, A. and Bekaert, G. (2002b). Regime switches in interest rates. *Journal of Business & Economic Statistics*, 20(2):163–182.

Ang, A. and Bekaert, G. (2004). How regimes affect asset allocation. *Financial Analysts Journal*, 60(2):86–99.

Ang, A. and Bekaert, G. (2007). Stock return predictability: Is it there? *The Review of Financial Studies*, 20(3):651–707.

Ang, A. and Chen, J. (2002). Asymmetric correlations of equity portfolios. *Journal of Financial Economics*, 63(3):443–494.

Ang, A. and Timmermann, A. (2012). Regime changes and financial markets. *Annual Review of Financial Economics*, 4(1):313–337.

Asness, C. S., Frazzini, A., and Pedersen, L. H. (2019). Quality minus junk. *Review of Accounting Studies*, 24(1):34–112.

Baker, M. and Wurgler, J. (2006). Investor sentiment and the cross-section of stock returns. *The Journal of Finance*, 61(4):1645–1680.

Bender, J., Sun, X., Thomas, R., and Zdorovtsov, V. (2018). The promises and pitfalls of factor timing. *The Journal of Portfolio Management*, 44(4):79–92.

Breusch, T. S. and Pagan, A. R. (1980). The lagrange multiplier test and its applications to model specification in econometrics. *The Review of Economic Studies*, 47(1):239–253.

Campbell, J. Y. and Shiller, R. J. (1988). The dividend-price ratio and expectations of future dividends and discount factors. *The Review of Financial Studies*, 1(3):195–228.

Chen, S.-S. (2009). Predicting the bear stock market: Macroeconomic variables as leading indicators. *Journal of Banking & Finance*, 33(2):211–223.

Davidson, R. and MacKinnon, J. G. (1983). Small sample properties of alternative forms of the lagrange multiplier test. *Economics Letters*, 12(3-4):269–275.

Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22.

Diebold, F., Lee, J.-H., and Weinbach, G. (1994). Regime switching with time-varying transition probabilities. In Hargraves, C., editor, *Non-stationary Time Series Analysis and Cointegration*, chapter 10, pages 283–302. Oxford University Press.

Fama, E. F. and French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56.

Friedman, J., Hastie, T., and Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1.

Gray, S. F. (1996). Modeling the conditional distribution of interest rates as a regime-switching process. *Journal of Financial Economics*, 42(1):27–62.

Guidolin, M. (2009). Detecting and exploiting regime switching ARCH dynamics in US stock and bond returns. In Gregoriou, G., editor, *Stock Market Volatility*, chapter 5, pages 92 – 132. CRC press.

Guidolin, M. (2011). Markov switching models in empirical finance. In Drukker, D. M., editor, *Missing data methods: Time-series methods and applications (Advances in Econometrics, Vol. 27 Part 2)*, pages 1–86. Emerald Group Publishing Limited.

Guidolin, M. and Ono, S. (2006). Are the dynamic linkages between the macroeconomy and asset prices time-varying? *Journal of Economics and Business*, 58(5-6):480–518.

Guidolin, M. and Timmermann, A. (2007). Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control*, 31(11):3503–3544.

Guidolin, M. and Timmermann, A. (2008). Size and value anomalies under regime shifts. *Journal of Financial Econometrics*, 6(1):1–48.

Gupta, T. and Kelly, B. (2019). Factor momentum everywhere. *The Journal of Portfolio Management*, 45(3):13–36.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2):357–384.

Hamilton, J. D. (1990). Analysis of time series subject to changes in regime. *Journal of Econometrics*, 45(1-2):39–70.

Hastie, T., Tibshirani, R., and Friedman, J. (2017). *The elements of statistical learning: Data mining, inference, and prediction.* Springer, second edition.

Kim, C.-J. (1994). Dynamic linear models with markov-switching. *Journal of Econometrics*, 60(1-2):1–22.

Kole, E. (2019). Markov switching models: An example for a stock market index. *Available at SSRN.*

Kole, E. and Van Dijk, D. (2017). How to identify and forecast bull and bear markets? *Journal of Applied Econometrics*, 32(1):120–139.

Lange, R. J. (2020). Lecture notes in advanced time series econometrics.

Ledoit, O. and Wolf, M. (2008). Robust performance hypothesis testing with the Sharpe ratio. *Journal of Empirical Finance*, 15(5):850–859.

Nyberg, H. (2013). Predicting bear and bull stock markets with dynamic binary time series models. *Journal of Banking & Finance*, 37(9):3351–3363.

Schaller, H. and Norden, S. V. (1997). Regime switching in stock market returns. *Applied Financial Economics*, 7(2):177–191.

Shapiro, S. S. and Wilk, M. B. (1965). An analysis of variance test for normality (complete samples). *Biometrika*, 52(3/4):591–611.

Welch, I. and Goyal, A. (2008). A comprehensive look at the empirical performance of equity premium prediction. *The Review of Financial Studies*, 21(4):1455–1508.

Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: series B (statistical methodology)*, 67(2):301–320.

Appendix

A Derivation of the (Conditional) Score

The relation of the conditional score with the full score is:

$$U_t(\boldsymbol{\theta}) = \frac{\partial \ell\left(\mathcal{I}_t; \boldsymbol{\theta}\right)}{\partial \boldsymbol{\theta}} - \frac{\partial \ell\left(\mathcal{I}_{t-1}; \boldsymbol{\theta}\right)}{\partial \boldsymbol{\theta}}. \tag{37}$$

Kole (2019) shows that the score for the set of parameters $\boldsymbol{\lambda} = \{\boldsymbol{A}_k, \boldsymbol{\Sigma}_k : k \in 1, \cdots, K\}$ is:

$$U(\boldsymbol{\lambda}) = \frac{\partial \ell\left(\mathcal{I}_T; \boldsymbol{\lambda}\right)}{\partial \boldsymbol{\lambda}} = \sum_{t=1}^{T} \frac{\partial \log[\boldsymbol{p}(\boldsymbol{r}_t)']}{\partial \boldsymbol{\lambda}} \hat{\boldsymbol{S}}_{t|T}, \tag{38}$$

where the elements of $\log[\boldsymbol{p}(\boldsymbol{r}_t)]$ are given by:

$$\log[p(\boldsymbol{r}_t | S_t = k)] = \frac{1}{2}\left(n \log|2\pi| + \log(|\boldsymbol{\Sigma}_k^{-1}|) - (\boldsymbol{r}_t - \mu_{k,t})' \boldsymbol{\Sigma}_k^{-1} (\boldsymbol{r}_t - \mu_{k,t})\right) \quad \text{for } k = 1, \cdots, K. \tag{39}$$

Its partial derivatives with respect to $\boldsymbol{\mu}_{k,t}$ and $\boldsymbol{\Sigma}_k$ are:

$$\begin{aligned} \frac{\partial \log[p(\boldsymbol{r}_t | S_t = k)]}{\partial \boldsymbol{\mu}_{k,t}} &= -\frac{1}{2}\left(\frac{\partial (\boldsymbol{r}_t - \boldsymbol{\mu}_{k,t})' \boldsymbol{\Sigma}_k^{-1} (\boldsymbol{r}_t - \boldsymbol{\mu}_{\boldsymbol{k},\boldsymbol{t}})}{\partial \boldsymbol{\mu}_{k,t}}\right) \\ &= -\frac{1}{2}\left(-2\boldsymbol{\Sigma}_k^{-1}(\boldsymbol{r}_t - \boldsymbol{\mu}_{k,t})\right) \\ &= \boldsymbol{\Sigma}_k^{-1}(\boldsymbol{r}_t - \boldsymbol{\mu}_{k,t}) \quad \text{for } k = 1, \cdots, K, \end{aligned} \tag{40}$$

and:

$$\begin{aligned} \frac{\partial \log[p(\boldsymbol{r}_t | S_t = k)]}{\partial \boldsymbol{\Sigma}_k} &= -\frac{\partial}{\partial \boldsymbol{\Sigma}_k} \frac{1}{2}\left[n \log|2\pi| + \log|\boldsymbol{\Sigma}_k| + (\boldsymbol{r}_t - \boldsymbol{\mu}_{k,t})' \boldsymbol{\Sigma}_k^{-1} (\boldsymbol{r}_t - \boldsymbol{\mu}_{k,t})\right] \\ &= -\frac{\partial}{\partial \boldsymbol{\Sigma}_k} \frac{1}{2}\left[\log|\boldsymbol{\Sigma}_k| + \text{trace}\left(\boldsymbol{\Sigma}_k^{-1}(\boldsymbol{r}_t - \boldsymbol{\mu}_{k,t})(\boldsymbol{r}_t - \boldsymbol{\mu}_{k,t})'\right)\right] \\ &= -\frac{1}{2}\left[2\boldsymbol{\Sigma}_k^{-1} - \text{diag}\left(\boldsymbol{\Sigma}_k^{-1}\right) - 2\boldsymbol{\Sigma}_k^{-1}(\boldsymbol{r}_t - \boldsymbol{\mu}_{k,t})(\boldsymbol{r}_t - \boldsymbol{\mu}_{k,t})' \boldsymbol{\Sigma}_k^{-1} \right. \\ &\quad \left. + \text{diag}\left(\boldsymbol{\Sigma}_k^{-1}(\boldsymbol{r}_t - \boldsymbol{\mu}_{k,t})(\boldsymbol{r}_t - \boldsymbol{\mu}_{k,t})' \boldsymbol{\Sigma}_k^{-1}\right)\right], \quad \text{for } k = 1, \cdots, K. \end{aligned} \tag{41}$$

From equations (37) and (38) it follows that the conditional score for parameter set $\boldsymbol{\lambda}$ is:

$$U_t(\boldsymbol{\lambda}) = \frac{\partial \log[\boldsymbol{p}(\boldsymbol{r}_t)']}{\partial \boldsymbol{\lambda}} \hat{\boldsymbol{S}}_{t|t} + \sum_{\tau=1}^{t-1} \frac{\partial \log \boldsymbol{f}_\tau'}{\partial \boldsymbol{\lambda}} \left(\hat{\boldsymbol{S}}_{\tau|t} - \hat{\boldsymbol{S}}_{\tau|t-1}\right), \tag{42}$$

with the relevant partial derivatives as in equations (40) and (41).

Kole (2019) also shows that for the set of parameters $\boldsymbol{\beta} = \{\beta_{kj}, \rho_k : k, j \in 1, \cdots, K\}$, the score is given by:

$$
\begin{aligned}
U(\boldsymbol{\beta}) =& \frac{\partial \ell\left(\mathcal{I}_T;\boldsymbol{\beta}\right)}{\partial \boldsymbol{\beta}} \\
=& \sum_{t=2}^{T}\sum_{j=1}^{K}\sum_{k=1}^{K} \frac{\partial \log p_{kj,t}}{\partial \boldsymbol{\beta}} \Pr\left[S_t = k, S_{t-1} = j \mid \mathcal{I}_T\right] \\
& + \sum_{k=1}^{K} \frac{\partial \log \rho_k}{\partial \boldsymbol{\beta}} \Pr\left[S_1 = k \mid \mathcal{I}_T\right]
\end{aligned}
\tag{43}
$$

where the partial derivatives are:

$$
\frac{\partial \log p_{kj,t}}{\partial \beta_{ij}} = \begin{cases} (1 - p_{kj,t})\boldsymbol{x}_{t-1} & \text{if } k = i \\ -p_{ij,t}\boldsymbol{x}_{t-1} & \text{if } k \neq i, \end{cases}
\tag{44}
$$

and $\Pr\left[S_t = k, S_{t-1} = j \mid \mathcal{I}_T,\right]$ and $\Pr\left[S_1 = k \mid \mathcal{I}_T\right]$ are calculated using the recursions in section 3.2.1. From equations (37) and (43), the conditional score for parameter set $\boldsymbol{\beta}$ is:

$$
\begin{aligned}
U_t(\boldsymbol{\beta}) =& \frac{\partial \ell_t\left(y_t \mid \mathcal{I}_{t-1};\boldsymbol{\beta}\right)}{\partial \boldsymbol{\beta}} \\
=& \sum_{j=1}^{K}\sum_{k=1}^{K} \frac{\partial \log p_{kj,t}}{\partial \boldsymbol{\beta}} \Pr\left[S_t = k, S_{t-1} = j \mid \mathcal{I}_t,\right] + \\
& \sum_{\tau=2}^{t-1}\sum_{j=1}^{K}\sum_{k=1}^{K} \frac{\partial \log p_{kj,\tau}}{\partial \boldsymbol{\beta}} \left(\Pr\left[S_\tau = k, S_{\tau-1} = j \mid \mathcal{I}_t\right] - \Pr\left[S_\tau = k, S_{\tau-1} = j \mid \mathcal{I}_{t-1}, \boldsymbol{\beta}\right]\right) + \\
& \sum_{k=1}^{K} \frac{\partial \log \rho_k}{\partial \boldsymbol{\beta}} \left(\Pr\left[S_1 = k \mid \mathcal{I}_t\right] - \Pr\left[S_1 \mid \mathcal{I}_{t-1}, \boldsymbol{\beta}\right]\right),
\end{aligned}
\tag{45}
$$

with the relevant partial derivatives as in equation (44).

B Descriptive Statistics

Table 9: Descriptive statistics of the explanatory variables.

Variable	Mean	sd	min	max
vxo	20.19	9.02	7.54	113.03
TED rate	0.52	0.41	0.09	3.97
yield spread	1.60	1.09	-0.77	3.81
prior r QMJ	0.54	3.21	-14.85	16.43
prior r SMB	0.06	3.59	-24.92	25.93
prior r HML	0.05	3.47	-15.64	19.97
prior r UMD	0.78	5.52	-41.11	21.94
prior r MKT	1.38	5.63	-32.45	29.33
MA 2500 QMJ	13.57	13.65	-10.62	79.20
MA 2500 HML	-14.48	14.89	-49.56	14.91
MA 2500 SMB	-5.07	15.64	-51.30	27.46
MA 2500 UMD	19.98	29.61	-45.17	106.36
MA 2500 MKT	43.58	30.43	-41.21	141.11

Table 10: A correlation matrix of the explanatory variables.

	vxo	TED rate	yield spread	prior r QMJ	prior r SMB	prior r HML	prior r UMD	prior r MKT	MA 2500 QMJ	MA 2500 HML	MA 2500 SMB	MA 2500 UMD	MA 2500 MKT
vxo	1.00												
TED rate	0.47	1.00											
yield spread	0.07	-0.14	1.00										
prior r QMJ	0.31	0.20	-0.05	1.00									
prior r SMB	-0.15	-0.15	0.13	-0.53	1.00								
prior r HML	-0.07	-0.06	0.11	-0.08	-0.14	1.00							
prior r UMD	-0.03	0.04	-0.07	0.36	-0.04	-0.25	1.00						
prior r MKT	-0.45	-0.21	-0.06	-0.62	0.35	-0.13	-0.28	1.00					
MA 2500 QMJ	0.31	-0.06	0.13	-0.04	0.13	0.08	-0.08	-0.03	1.00				
MA 2500 HML	-0.13	-0.29	0.15	-0.01	-0.01	0.03	-0.08	-0.06	-0.20	1.00			
MA 2500 SMB	-0.01	0.28	0.10	0.04	-0.04	0.15	0.04	-0.05	0.01	0.20	1.00		
MA 2500 UMD	0.15	0.16	-0.21	0.03	0.01	0.14	0.04	-0.02	0.43	-0.52	0.10	1.00	
MA 2500 MKT	-0.04	0.20	-0.42	0.07	-0.12	-0.07	0.16	0.06	-0.20	-0.60	-0.28	0.40	1.00

C Monte Carlo Algorithm for Multi-Period Returns

Algorithm 1: Monte Carlo simulation for total h period return and variance estimation at time t

Input: Estimates for the state beliefs $\hat{\boldsymbol{S}}_{t|t}$, regime-dependent means $\hat{\boldsymbol{\mu}}_k$ and covariance matrix $\hat{\boldsymbol{\Sigma}}_k$, the estimated transition probability matrix $\mathbf{P}$ and return vector at time t: $\boldsymbol{r}_t$

Number of `paths` $= N$;

Number of `periods` $= H$;

Calculate $\hat{\boldsymbol{S}}_{t+1|t}$ using $\hat{\boldsymbol{S}}_{t+1|t} = \mathbf{P}\hat{\boldsymbol{S}}_{t|t}$;

for *n in 1:N* **do**

  Draw H i.i.d state innovations $(\boldsymbol{\varepsilon}_{(1),k}, \cdots, \boldsymbol{\varepsilon}_{(H),k}) \quad \forall \quad k \in 1, \cdots, K$ from the multivariate normal distribution $\mathrm{N}(0, \hat{\boldsymbol{\Sigma}}_k)$;

  for *h in 1:H* **do**

    Draw a state realization k^* from a categorical distribution with probabilities $\hat{\boldsymbol{S}}_{t+h|t+h-1}$;

    Construct the return vector for path n at time $t+h$ as:

    $\boldsymbol{r}_{t+h}^{(n)} = \hat{\boldsymbol{\mu}}_{k^*} + \boldsymbol{\varepsilon}_{(h),k^*}$;

    Calculate $\hat{\boldsymbol{S}}_{t+h+1|t+h}$ using $\hat{\boldsymbol{S}}_{t+h+1|t+h} = \mathbf{P}\boldsymbol{S}_{t+h}$, with $\boldsymbol{S}_{t+h}$ a vector of 0's with a 1 at position k^*;

  end

  Calculate the total return over h periods in path n:

  $\boldsymbol{r}_{t:h}^{(n)} = \prod_{h=1}^{H}(1 + \boldsymbol{r}_{t+h}^{(n)}/100)$

end

Calculate the expected total return over h periods as:

$\hat{\boldsymbol{r}}_{t:h|t} = \frac{1}{N}\sum_{n=1}^{N} r_{t:h}^{(n)}$;

and the covariance matrix of h-period returns as

$\hat{\boldsymbol{\Sigma}}_{t:h|t} = \frac{1}{N-1}\sum_{n=1}^{N}\left(\boldsymbol{r}_{t:h}^{(n)} - \hat{\boldsymbol{r}}_{t:h|t}\right)\left(\boldsymbol{r}_{t:h}^{(n)} - \hat{\boldsymbol{r}}_{t:h|t}\right)'$;

D Variable Selection Procedures

D.1 Iterations in Step-Wise Score Test

Table 11: The selected variable/ transition combinations with the corresponding log-likelihood from each iteration of the step-wise procedure using the score test, as outlined in section 3.4.2

Iteration	*Variable*	*Transition*	*Log-likelihood*	*Iteration*	*Variable*	*Transition*	*Log-likelihood*
1	vxo	2 -> 3	14322.0	12	yield spread	1 -> 2	14274.9
2	vxo	3 -> 3	14316.5	13	vxo	2 -> 2	14267.4
3	MA 2500 SMB	3 -> 3	14310.1	14	MA 2500 HML	1 -> 2	14260.7
4	prior r SMB	2 -> 2	14305.8	15	prior r HML	3 -> 2	14259.2
5	MA 2500 SMB	3 -> 2	14301.9	16	MA 2500 QMJ	3 -> 2	14256.4
6	prior r QMJ	1 -> 2	14299.1	17	MA 2500 QMJ	3 -> 3	14252.6
7	MA 2500 MKT	1 -> 3	14295.6	18	MA 2500 QMJ	1 -> 3	14251.2
8	prior r UMD	1 -> 2	14292.5	19	yield spread	1 -> 3	14247.5
9	MA 2500 SMB	2 -> 3	14289.4	20	MA 2500 SMB	1 -> 3	14245.1
10	vxo	1 -> 3	14287.9	21	prior r HML	1 -> 3	14241.6
11	MA 2500 SMB	1 -> 2	14282.9				

D.2 Choice of λ

Table 12: Model evaluation metrics for each set of λ's that has been tried as the regularization parameters in the LASSO estimation.

	choice of λ			*Measure*						
λ set	λ 1	λ 2	λ 3	Neg. loglik.	AIC	BIC	Dens. Spec. Test. (p-val)	# nonzero param. 1	# nonzero param. 2	# nonzero param. 3
1	0	0	0	14101	28730	30185	1.622e-06	13	13	13
2	0.032	0.063	0.066	14161	28830	30230	1.777e-07	13	11	13
3	0.04	0.076	0.092	14179	28847	30191	2.594e-07	11	11	13
4	0.05	0.091	0.127	14193	28873	30218	3.214e-07	11	11	13
5	0.064	0.11	0.176	14225	28858	29982	4.732e-07	8	9	10
6	0.08	0.132	0.244	14268	28871	29830	3.981e-07	8	6	7
7	0.101	0.159	0.338	14253	28784	29549	5.748e-05	4	4	6
8	0.128	0.192	0.468	14281	28801	29457	0.0001342	4	3	3
9	0.161	0.231	0.648	14311	28820	29366	0.0002854	1	3	2
10	0.203	0.279	0.898	14332	28812	29220	0.000485	1	0	0

E Coefficients of Three State Model with Explanatory Variables

E.1 Coefficients of 3TVTP Model

Table 13: The coefficients for the parameters explaining the transition probabilities in the 3TVTP model, estimated on the full sample, and using standardized explanatory variables.

from	1 *(low-volatility)*		2 *(medium-volatility)*		3 *(high-volatility)*	
to	2	3	2	3	2	3
constant	-3.61^{***}	-8.84^{***}	3.12^{***}	0.5	1.23^{*}	1.51^{**}
vxo		1.3^{**}	2.53^{***}	3.86^{***}		0.94^{***}
	(-0.1)	(1.9)	(12.2)	(10.5)	(-12.5)	(17.8)
TED rate	-	-	-	-	-	-
yield spread	-1.94^{***}	2.15^{***}	-	-	-	-
	(-10.1)	(3.3)				
prior r QMJ	1.53^{***}		-	-	-	-
	(7.9)	(-0.1)				
prior r SMB	-	-	-0.41^{***}		-	-
			(-6.3)	(2.8)		
prior r HML		1.2^{**}	-	-	-0.29^{*}	
	(-0.1)	(1.8)			(-5.3)	(3.9)
prior r UMD	1.56^{*}		-	-	-	-
	(8.1)	(-0.1)				
prior r MKT	-	-	-	-	-	-
MA 2500 QMJ		-1.63^{**}	-	-	1.56^{**}	1.17^{*}
	(0.1)	(-2.4)			(12.8)	(1.4)
MA 2500 HML	1.47^{***}		-	-	-	-
	(7.6)	(-0.1)				
MA 2500 SMB	-3.04^{***}	2.25^{*}		-0.57^{***}	-2.16^{**}	-2.36^{**}
	(-15.8)	(3.5)	(3.9)	(-4.1)	(-7.8)	(-16.2)
MA 2500 UMD	-	-	-	-	-	-
MA 2500 MKT		-1.64^{***}	-	-	-	-
	(0.1)	(-2.4)				

Note: Marginal effects are shown in parentheses and are denoted in %. ***,**,* display significance at the 10%, 5%, and 1% level respectively, for a t-test testing the null that the coefficient is 0. Standard errors are estimated with the OPG method. Arrival state 1 is not shown since it is the reference category. The marginal effect that a variable has on transitions to regime 1 is calculated as 0 minus the sum of the marginal effects of transitions to regime 2 and 3.

E.2 Coefficients of 3CE Model

Table 14: The coefficients for matrix $\boldsymbol{A}_k$ in a 3-regime model with explanatory variables (3CE) estimated using the group-LASSO with regularization parameter λ equal to 0.101, 0.159, and 0.338 in regimes 1, 2, 3 respectively, using the full sample.

Variable	Regime 1 *(low-volatility)*					Regime 2 *(medium-volatility)*					Regime 3 *(high-volatility)*				
						Coefficient (× 100)									
	QMJ	SMB	HML	UMD	MKT	QMJ	SMB	HML	UMD	MKT	QMJ	SMB	HML	UMD	MKT
Constant	1.943	0.725	6.311	10.462	17.168	10.127	-15.209	-28.684	42.837	31.942	39.733	20.038	98.268	-47.635	-138.838
vxo	-0.078	0.174	-0.001	0.567	1.188	-	-	-	-	-	-0.961	-0.317	-2.534	-1.32	3.96
TED rate	-	-	-	-	-	-	-	-	-	-	9.465	-1.360	-1.581	8.191	-34.336
yield spread	-	-	-	-	-	2.35	7.559	7.525	-1.511	-8.281	-	-	-	-	-
prior r QMJ	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
prior r SMB	-	-	-	-	-	-	-	-	-	-	-0.392	0.114	-0.619	2.57	1.272
prior r HML	0.009	-1.196	1.585	-0.376	-1.522	-	-	-	-	-	-	-	-	-	-
prior r UMD	-	-	-	-	-	1	-0.795	-0.168	1.415	-1.582	-	-	-	-	-
prior r MKT	-0.382	1.597	0.536	-1.483	-2.535	-0.518	1.007	-0.133	-0.295	-0.255	-3.535	2.683	0.514	-5.277	0.968
MA 2500 QMJ	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
MA 2500 HML	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
MA 2500 SMB	-	-	-	-	-	0.063	-0.017	0.07	0.025	-0.093	0.237	-0.225	0.191	-0.607	-0.862
MA 2500 UMD	-	-	-	-	-	-	-	-	-	-	0.001	0.015	0.011	0.057	0.004
MA 2500 MKT	0.111	-0.157	-0.041	0.048	-0.044	-	-	-	-	-	-	-	-	-	-

F Additional Figures

F.1 In-Sample Smoothed State Beliefs for three state models

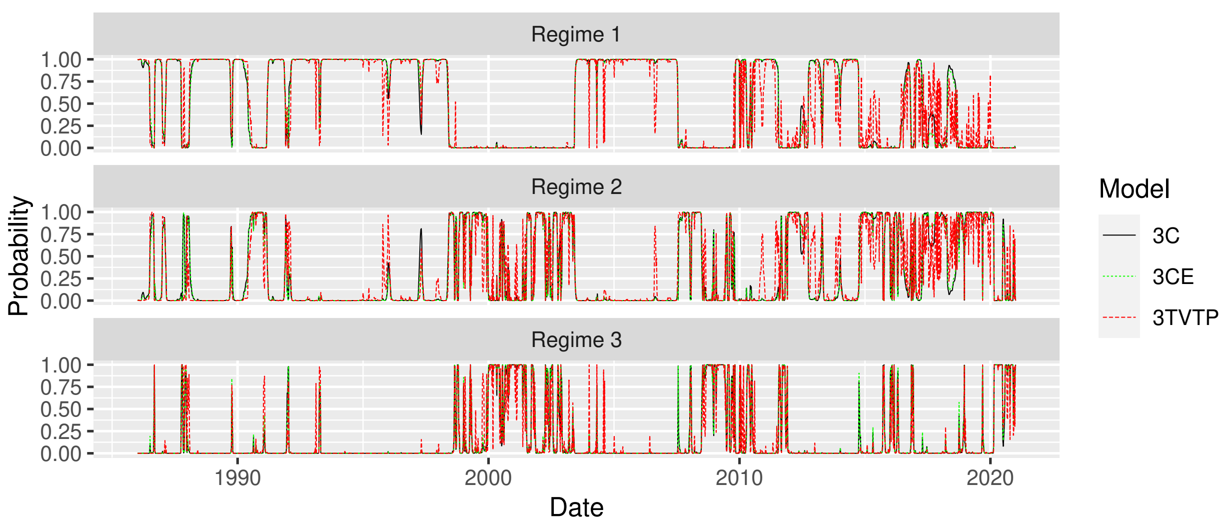

Figure 11: The smoothed state probabilities estimated on the full sample for the 3C, 3TVTP, and 3CE model.

F.2 In-Sample Cumulative Returns

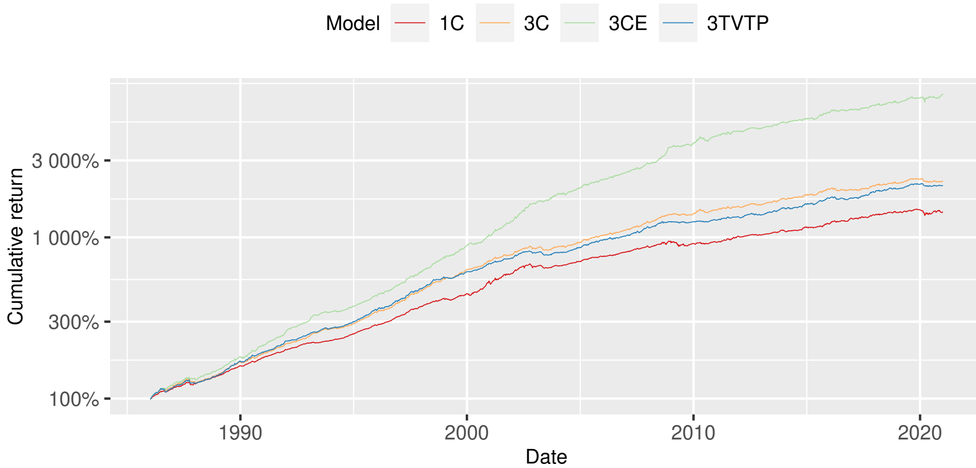

Figure 12: The cumulative returns from the portfolios formed by solving the portfolio optimization problem in equation (24) for 1-period ahead returns using inputs of the 3C, 3TVTP, 3CE, and 1C models, estimated using the full sample. Moments of the Markov-switching models for input in equation (24) are calculated using predicted state beliefs ($\hat{\boldsymbol{S}}_{t+1|t}$) as in equations (22, 23). The model which does not take regimes into account (1C) takes the historical mean and covariance of h-period returns of the full sample. The Y-axis is transformed to the log-10 scale.

F.3 Mean Returns Differentiated by Number of Correctly Estimated Factor Return Directions

Table 15: Mean returns per correctly predicted factor return direction

	correctly predicted factor return direction				
model	1	2	3	4	5
4C	-0.28	-0.01	0.25	0.47	0.71
3C	-0.30	-0.03	0.26	0.49	0.82
2C	-0.20	0.01	0.22	0.48	0.84
1C	-0.23	-0.03	0.24	0.42	0.58
3CE	-0.34	0.01	0.27	0.59	0.89
3TVTP	-0.30	-0.03	0.26	0.49	0.74

F.4 OOS Weights for Multi-Period Investment Horizons

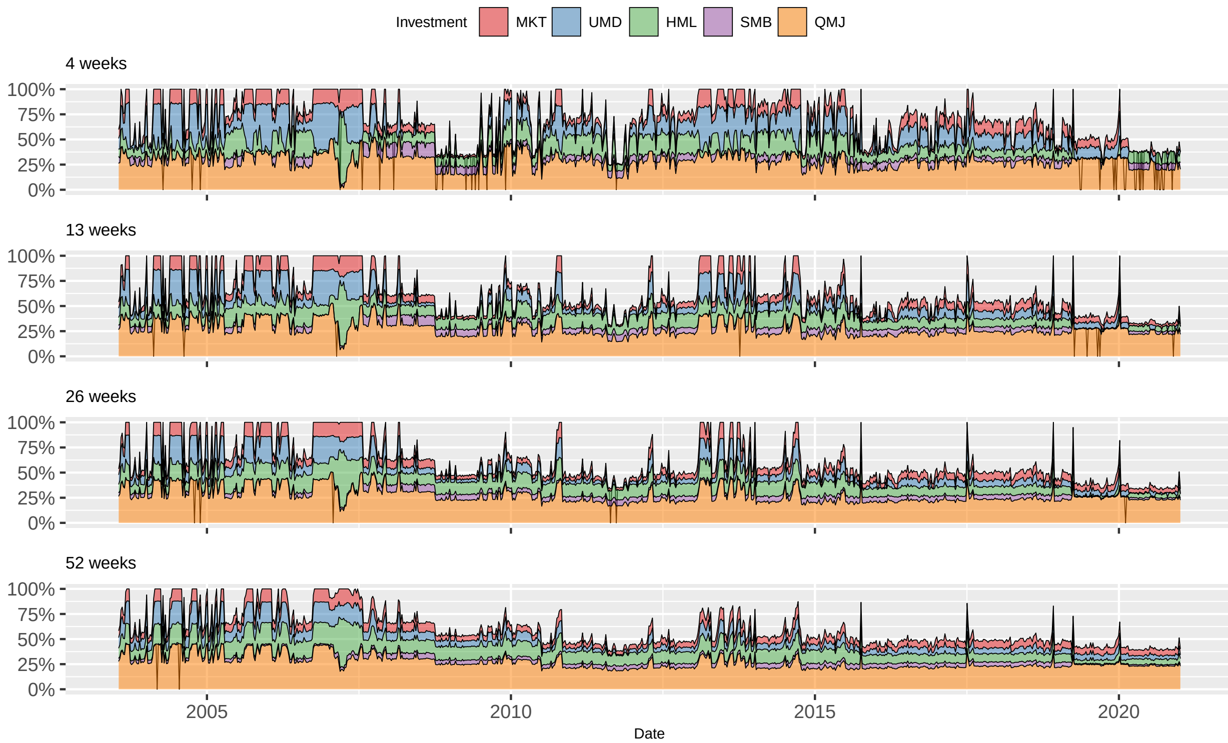

Figure 13: The OOS weights from the optimal solution of the portfolio optimization problem in equation (24) for 4-, 13-, 26-, and 52-week returns at each time t. $1 - w_{MKT}$ is invested in the risk-free asset, following from (24). Moments are constructed with regime-dependent return vectors and state covariance matrices estimated with and expanding window from a three-regime model with constant transition probabilities (3C). Moments are calculated using algorithm 1.

G Robustness Checks

G.1 Estimation of Five Single-Factor Three-Regime Models

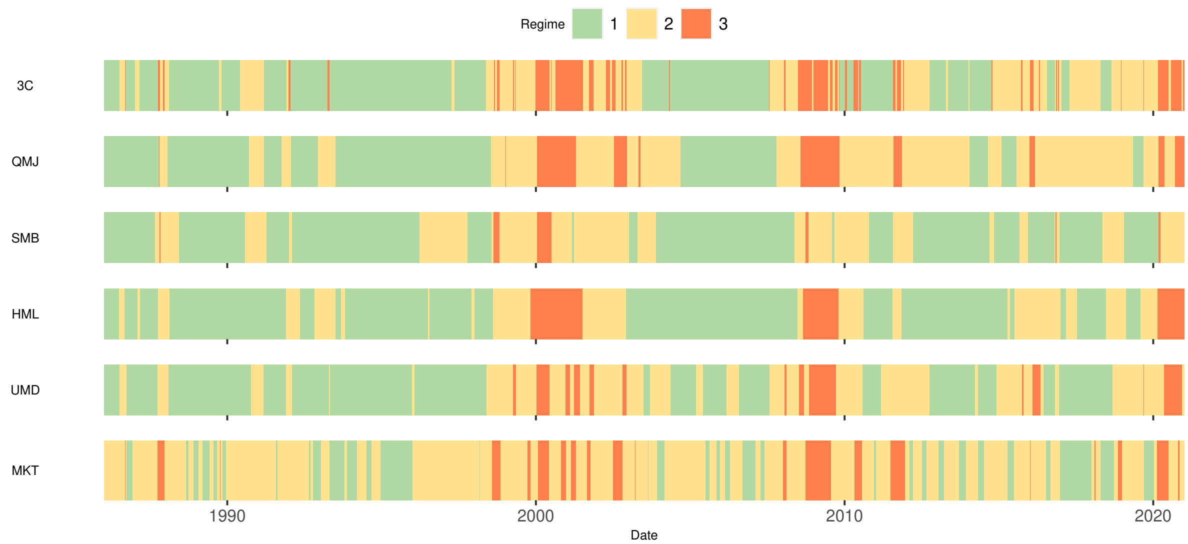

Figure 14: The sequences of most likely smoothed state beliefs of the 3C model and five three-regime models that model the process of each factor separately.

Table 16: The mean and variance of five three-regime models that model the process of each factor separately

	Mean					Variance				
	QMJ	SMB	HML	UMD	MKT	QMJ	SMB	HML	UMD	MKT
Regime 1	0.094	-0.051	-0.014	0.209	0.465	0.285	0.897	0.530	0.734	0.906
Regime 2	0.064	0.119	0.021	0.308	0.189	1.055	2.154	1.828	3.423	3.733
Regime 3	0.191	-0.323	0.062	-1.034	-0.503	5.376	14.721	6.469	23.654	22.307

G.2 Portfolio Optimization Problem with Different Constraints

Table 17: OOS performance measures of the optimal portfolio solution from equation (24) for one-period ahead returns, with the no-short-selling constraint substituted by the constraint :$-1 \leq w_{i,t}$, using inputs from the 3C, 3TVTP, 3CE, and 1C model, estimated with the expanding window approach.

Model	*Performance measure*				*Sharpe ratio difference*		
	$\bar{r}^e$ (%)	σ (%)	Utility ($\times$ 10,000)	SR (annualized)	1C	3TVTP	3CE
3C	0.032	0.444	5.260	0.516	-0.263	-0.748^{***}	-0.079^{*}
3CE	0.082	0.994	9.039	0.595	-0.184	-0.669^{*}	
3TVTP	0.099	0.567	11.826	1.264	0.485^{*}		
1C	0.078	0.725	9.441	0.779			

Note: Moments of the Markov-switching models for input in equation (24) are calculated using predicted state beliefs ($\hat{\boldsymbol{S}}_{t+1|t}$) as in equations (22, 23). The return vector and covariance matrix of the model which does not take regimes into account (1C) are the historical mean and covariance up to each of the 10 splits of the expanding window. The weekly mean excess portfolio return, standard deviation, realized utility, and Sharpe ratios are reported. Weekly Sharpe ratios have been annualized by multiplying by $\sqrt{52}$. P-values for the test of differences in Sharpe ratios are in parentheses. They are calculated using the test from Ledoit and Wolf (2008). This test is conducted on the non-annualized Sharpe ratios. ***,** ,* display significance at the 10%, 5%, and 1% level respectively.

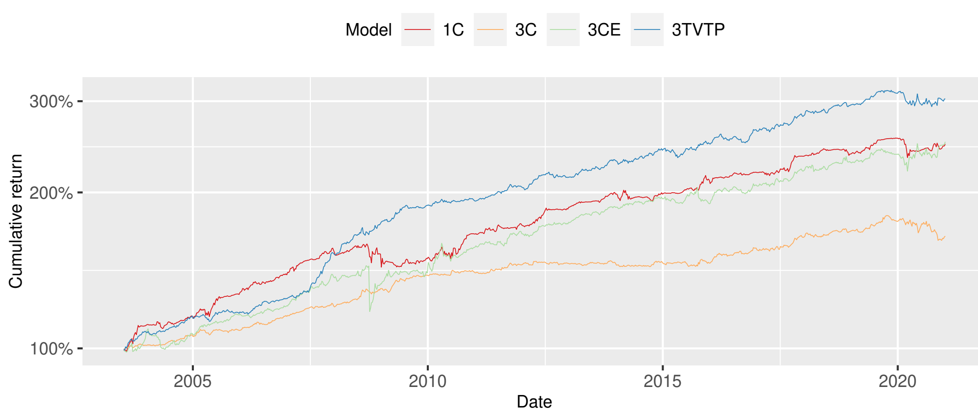

Figure 15: The cumulative returns from the portfolios formed by solving for the optimal portfolio solution from equation (24) for one-period ahead returns, with the no-short-selling constraint substituted by the constraint :$-1 \leq w_{i,t}$, using inputs of the 3C, 3TVTP, 3CE, and 1C models, estimated with the expanding window approach. Moments of the Markov-switching models for input in equation (24) are calculated using predicted state beliefs ($\hat{\boldsymbol{S}}_{t+1|t}$) as in equations (22, 23). The model which does not take regimes into account (1C) takes the historical mean and covariance of h-period returns up to each of the 10 splits of the expanding window as input for the portfolio optimization problem. The Y-axis is transformed to the log-10 scale.

Table 18: OOS performance measures of the optimal portfolio solution from equation (24) for h-period ahead returns, with the no-short-selling constraint substituted by the constraint :$-1 \leq w_{i,t}$, using inputs from the 3C and 1C model, estimated with the expanding window approach.

Model	*Performance measure*				
	$\bar{r}^e$ (%)	σ (%)	Utility (*10,000)	SR (annualized)	Sharpe ratio difference
			4 week horizon		
3C	0.447	1.758	53.468	0.916	0.504***
1C	0.178	1.563	10.168	0.412	
			26 week horizon		
3C	1.339	2.872	163.465	0.933	0.465***
1C	0.700	2.993	-476.386	0.468	
			26 week horizon		
3C	2.587	4.221	317.712	0.867	0.420***
1C	1.379	4.365	-5008.140	0.447	
			52 week horizon		
3C	4.755	6.243	593.006	0.762	0.310**
1C	2.662	5.896	-57452.506	0.452	

Note: Moments of the 3C model for input in equation (24) are calculated using algorithm 1. The return vector and covariance matrix of the model which does not take regimes into account (1C) are the historical mean and covariance of h-period returns up to each of the 10 splits of the expanding window. The h-period mean excess portfolio return, standard deviation, realized utility, and Sharpe ratios are reported. Weekly Sharpe ratios have been annualized by multiplying the 4-, 13- and 26-period returns by $\sqrt{13}, \sqrt{4}, \sqrt{2}$ respectively. P-values for the test of differences in Sharpe ratios are in parentheses. They are calculated using the test from Ledoit and Wolf (2008). This test is conducted on the non-annualized Sharpe ratios. ***,**,* display significance at the 10%, 5%, and 1% level respectively.

G.3 Consistency of Out-of-Sample Procedure

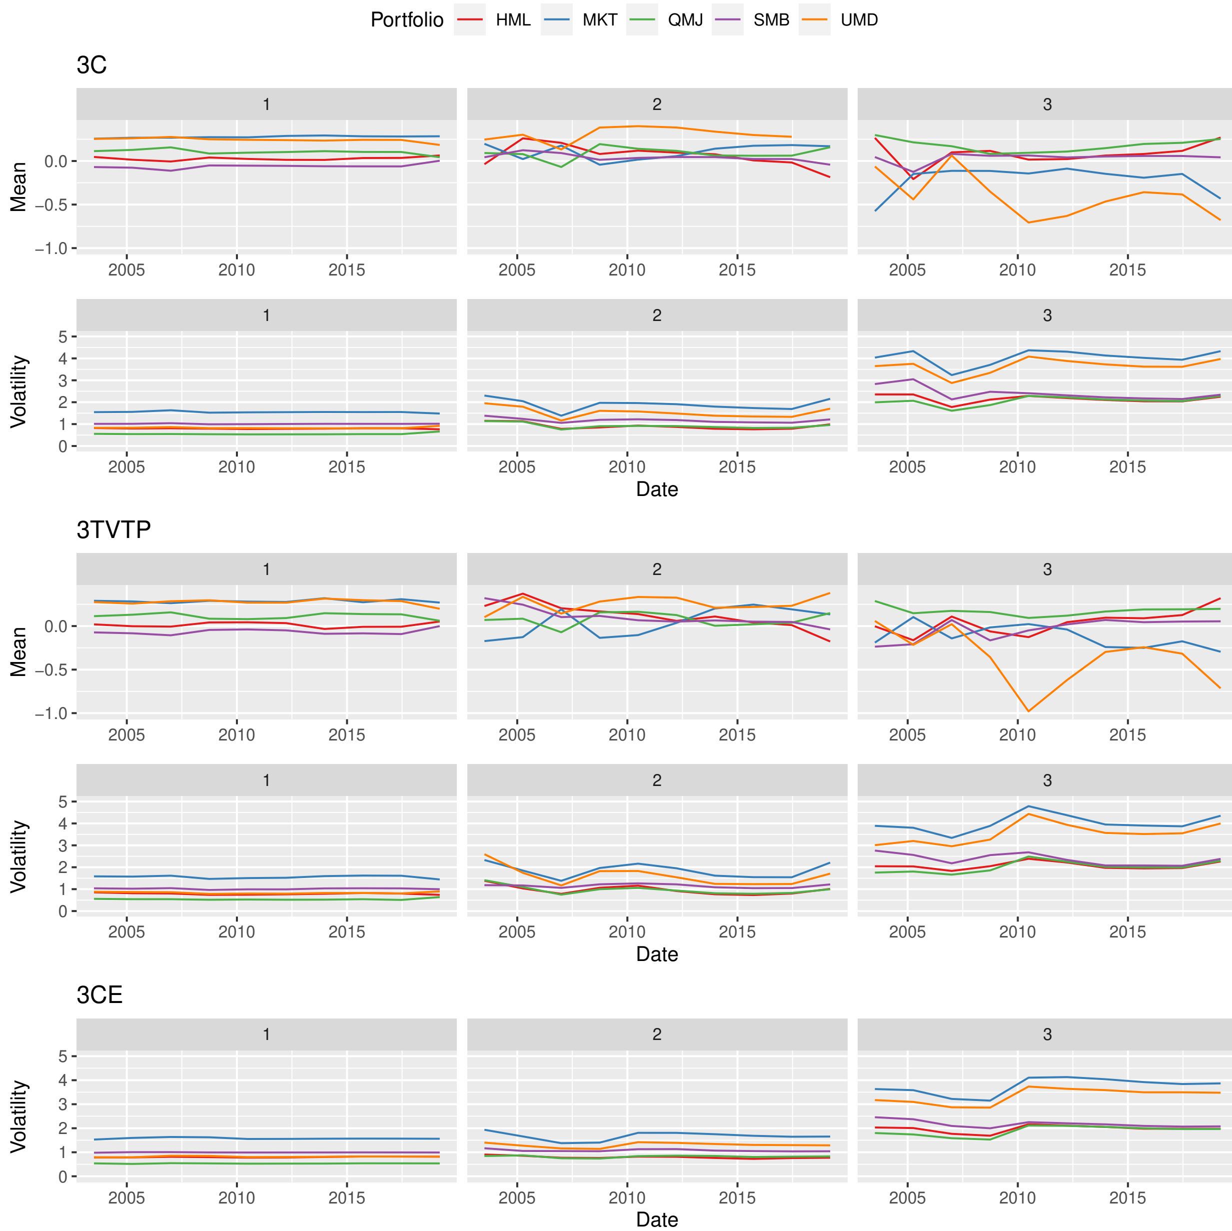

Figure 16: The values for the mean return and volatility of each of the portfolios after training the models with an expanding window method over each of the 10 splits. Mean returns of the 3CE model are not displayed since they are dependent on explanatory variables. Volatility of each of the factors is taken as the square root of the values on the diagonal of the estimated regime-dependent covariance matrices. It is displayed instead of the variance due to the scale of the Y-axis.

G.4 Consistency of Variable Selection Procedure

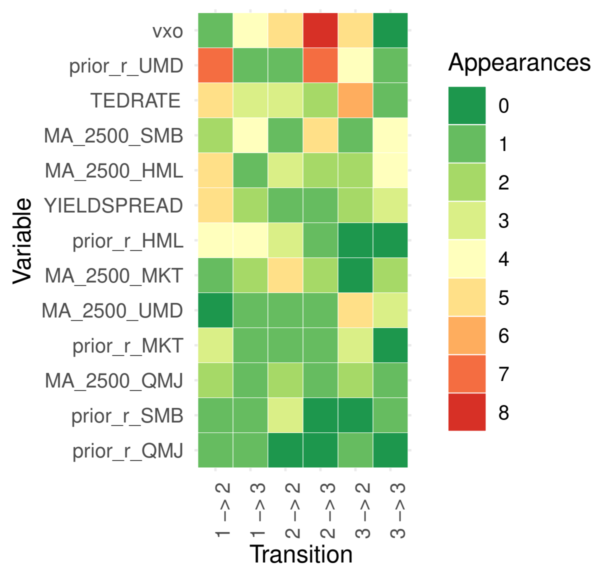

Figure 17: A heatmap with the number of times each variable is selected for each transition after training the models with an expanding window method over each of the 10 splits. Variable selection is performed using the stepwise procedure with the score test. The variables are ranked vertically from most (top) to least (bottom) selected.

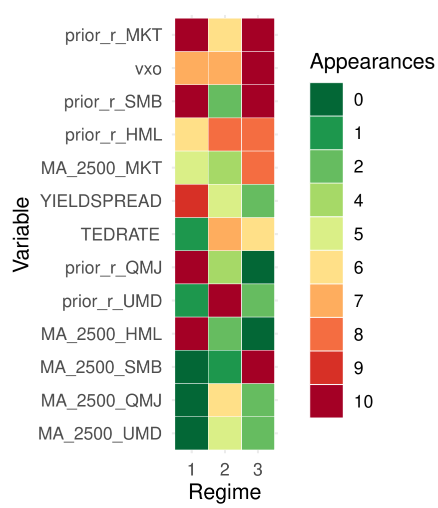

Figure 18: A heatmap with the number of times each variable is selected for the equation for returns in each state, after training the models with an expanding window method over each of the 10 splits. Variable selection is performed using the LASSO, with λ equal to 0.101, 0.159, and 0.338, in regimes 1, 2, 3 respectively (the λ's chosen over the full sample). Variables are ranked vertically from most (top) to least (bottom) selected.

H Programming code

All programming code is written in `R`, and can be found in the file `code_appendix.R`. Further elaboration of the code is also included in this file using comments.