



Munich Personal RePEc Archive

Markov-switching Asset Allocation: Do Profitable Strategies Exist?

Bulla, Jan and Mergner, Sascha and Bulla, Ingo and
Sesboüé, André and Chesneau, Christophe

Université de Caen, Département de Mathématiques, Quoniam
Asset Management GmbH, Georg-August-Universität Göttingen,
Institut für Mikrobiologie und Genetik

7 January 2010

Online at <https://mpra.ub.uni-muenchen.de/21154/>
MPRA Paper No. 21154, posted 07 Mar 2010 00:27 UTC

Markov-switching Asset Allocation: Do Profitable Strategies Exist?

Jan Bulla^{1*}, Sascha Mergner², Ingo Bulla³
André Sesboüé¹, Christophe Chesneau¹

¹Université de Caen

Département de Mathématiques

Bd Maréchal Juin, BP 5186, 14032 Caen Cedex, France

²Quoniam Asset Management GmbH

Westhafen Tower, Westhafenplatz 1

60327 Frankfurt am Main, Germany

³Institut für Mikrobiologie und Genetik

Abteilung für Bioinformatik

Georg-August-Universität Göttingen

Goldschmidtstr. 1, 37077 Göttingen, Germany

January 7th 2010

*Corresponding author, Phone: +33 (0) 250505364, Fax: +33 (0) 231567320 Email: bulla@math.unicaen.fr, jan.bulla@yahoo.com. The e-mail addresses and phone numbers of the other authors are Sascha.Mergner@quoniam.de (+49 (0) 69 74384 224), ibulla@uni-goettingen.de (+49 (0) 551 39 14627), sesboue@math.unicaen.fr (+33 (0) 231567465), and chesneau@math.unicaen.fr (+33 (0) 231567474)

Acknowledgments: We would like to sincerely thank P. Thomson, D. Vere-Jones, and W. Zucchini for their inspiring comments and support. We would also like to render thanks to the participants of the Cherry Bud Workshop 2007 and the 17th NZESG for helpful feedback. Not to forget, we thank Gisela and Klaus Bulla for editorial assistance. Jan Bulla was supported by a fellowship within the Postdoc-Program of the German Research Foundation (DFG).

Abstract

This paper proposes a straightforward Markov-switching asset allocation model, which reduces the market exposure to periods of high volatility. The main purpose of the study is to examine the performance of a regime-based asset allocation strategy under realistic assumptions, compared to a buy and hold strategy. An empirical study, utilizing daily return series of major equity indices in the US, Japan, and Germany over the last 40 years, investigates the performance of the model. In an out-of-sample context, the strategy proves profitable after taking transaction costs into account. For the regional markets under consideration, the volatility reduces on average by 41%. Additionally, annualized excess returns attain 18.5 to 201.6 basis points.

Keywords: Hidden Markov model, Markov-switching model, asset allocation, timing, volatility regimes, daily returns.

JEL classification codes: C13, C15, C22, E44, G11, G15.

Introduction

Asset allocation decisions represent the most important single determinant of an investor's performance (Brinson et al., 1991). In this paper, the effects and consequences of regime-switching on asset allocation are analyzed. Considering a simple two-asset world, we employ a Markov-switching approach in which the decision to invest in the stock market or in cash depends on the prevailing market regime. The main purpose of this study is to examine the profitability of a regime-based asset allocation strategy after taking transaction costs explicitly into account. The present paper contributes an empirical analysis of the in-sample and out-of-sample

performance of a Markov-switching asset allocation model applied to daily stock market return series in the US, Japan, and Germany over several decades.

Regime-switching models in finance

In the past two decades, regime-switching models have attracted increasing interest by researchers in the fields of macroeconomics and financial time series. Markov-switching models represent time series models with a latent variable component where an unobserved Markov process drives the observation-generating distribution. Despite the flexibility of the hidden Markov model (HMM) and its widespread use among engineers in the field of signal-processing, applications of Markov-switching models (a synonym for HMM) to economics and financial econometrics evolved mainly after the seminal work of Hamilton (1989).

The introduction of time-varying parameter models to the scientific community dates back to Quandt (1958) who presented an estimation approach for a linear regression system with two regimes. In a later study, Quandt (1972) refined his techniques and applied them to analyze disequilibria in the housing market. In the following year, Goldfeld and Quandt (1973) introduced Markov-switching regression. Hamilton (1989) focused on autoregressive models with Markov-switching parameters. During the same period, researchers in the field of speech recognition successfully worked on related models. These can be traced back to Baum and Petrie (1966) and Baum et al. (1970), who laid the groundwork for the influential works of Dempster et al. (1977) and Rabiner (1989).

The findings relevant to the subsequent analysis belong to the field of modeling daily return series with Markov-switching mixture distributions. The study of Turner et al. (1989) may be the first in this context, and others followed, e.g., Rydén et al. (1998); Linne (2002); Bialkowski (2003). For further details, we refer to the monographs of MacDonald and Zucchini (2009) and Cappé et al. (2007).

Asset allocation and regime-switching

The effects of regime-switching on asset allocation have been investigated using different approaches for different types of data. One of the earliest works on this subject has been presented by Ang and Bekaert (2002) who treated an asset allocation problem with shifting regimes from the perspective of a US investor. Their core analysis focused on a dataset of monthly Morgan Stanley Capital International (MSCI) total returns from January 1970 to December 1997. The authors' main conclusions were that the existence of a high-volatility bear market regime does not negate the benefits of international diversification and that the high volatility, high correlation regime tends to coincide with a bear market. In a subsequent study, Ang and Bekaert (2004) extended their data basis and changed the period under investigation. They found that the regime-switching strategy dominates static strategies out-of-sample for a global all-equities portfolio, and that the model proposes to switch primarily to cash in a persistent high-volatile market. Utilizing an autoregressive Markov-switching model, Graffund and Nilsson (2003) analyzed monthly returns for Germany/Japan (1950–1999) and US/UK (1900–1999). They investigated an intertemporal asset allocation problem for investors who dynamically rebalance their portfolio every month. The authors highlighted the economic importance of regimes and pointed out that optimal portfolio weights are clearly dependent on the current regime. Bauer et al. (2004) focused on monthly returns from January 1976 to December 2002 of a six-asset portfolio consisting of equities, bonds, commodities, and real estate. They observed changing correlations and volatilities among assets, and demonstrated a significant information gain by using a regime-switching instead of a standard mean-variance optimization strategy. The study of Guidolin and Timmermann (2005) analyzed the FTSE All Share stock market index, returns on 15-year government bonds and 1-month T-bills from January 1976 to December 2000. It presents strong evidence of regimes with different risk and return characteristics for UK stocks and bonds, and evidence of persistent bull and bear regimes for both series. Ammann and Verhofen (2006) estimated the four-factor model of Carhart (1997), using monthly data from January 1927

to December 2004. They found two clearly separable regimes with different mean returns, volatilities, and correlations. One of their key findings was that value stocks provide high returns in the high-variance state, whereas momentum stocks and the market portfolio perform better in the low-variance state. Finally, Hess (2006) examined the improvement of portfolio performance when imposing conditional CAPM strategies based on regime forecasts from an autoregressive Markov regime-switching behavior. Based on returns of the Swiss stock market and its 18 sectors from January 1973 to June 2001, his results indicate that regime switches are a valuable timing signal for portfolio rebalancing.

However, an important point with regard to the practical applicability of regime-switching models is the number of state changes, because frequent rebalancing of the portfolio is likely to eat up much of the potential excess returns (as described, e.g., in the studies of Bauer et al., 2004; Hess, 2006). Alternatively, transaction costs are often simply not explicitly taken into account (see e.g. Graflund and Nilsson, 2003; Ang and Bekaert, 2002, 2004; Guidolin and Timmermann, 2005), just as out-of-sample forecasts are not always included. This does not reduce the value of the descriptive capacities of the models, however, it limits their practical application.

The remainder of this article is organized as follows. Section 1 outlines the methodology; it introduces the basic concept of HMMs and the approach to regime-driven asset allocation. Section 2 gives a short description of the data and provide summary statistics. Section 3 discusses the empirical results, and Section 4 concludes.

1 Methodology

1.1 The basic hidden Markov model

The main characteristic of a HMM is a probability distribution of the observation X_t , $t = 1, \dots, T$ that is determined by the unobserved states S_t of a homogeneous and irreducible finite-state Markov chain. The switching behavior is governed by a *transition probability matrix* (TPM). Assuming a model with two states, the TPM

is of the form

$$\mathbf{\Pi} = \begin{pmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{pmatrix},$$

where p_{ij} , $i, j \in \{1, 2\}$ denote the probability of being in state j at time $t+1$ given a sojourn in state i at time t . The distribution of the observation at time t is specified by the *conditional* or *component distributions* $P(X_t = x_t | S_t = s_t)$. Assuming, for instance, a two-state model with Gaussian component distributions yields

$$x_t = \mu_{s_t} + \epsilon_{s_t}, \quad \epsilon_{s_t} \sim N(0, \sigma_{s_t}^2), \quad (1)$$

where $\mu_{s_t} \in \{\mu_1, \mu_2\}$ and $\sigma_{s_t}^2 \in \{\sigma_1^2, \sigma_2^2\}$. The parameters of a HMM are generally estimated by the method of maximum-likelihood. The likelihood function is available in a convenient form:

$$L(\theta) = \boldsymbol{\pi} \mathbf{P}(x_1) \mathbf{\Pi} \mathbf{P}(x_2) \mathbf{\Pi} \dots \mathbf{P}(x_{T-1}) \mathbf{\Pi} \mathbf{P}(x_T) \mathbf{1}',$$

where $\mathbf{P}(x_t)$ represents a diagonal matrix with the state-dependent conditional distributions as entries (MacDonald and Zucchini, 2009) and $\boldsymbol{\pi}$ denotes the initial distribution of the Markov chain. The estimation of the model parameters θ for stationary HMMs in this work follows Bulla and Berzel (2008).

1.2 Regime-switching asset allocation with daily return series

As summarized in Section , various studies provide insight into the effect of regime-switching on asset allocation. Although many different samples have been analyzed, previous work largely focuses on monthly returns. The results are encouraging, mostly finding strong evidence for different regimes with returns above or below the historical average. However, the literature so far has not yet considered daily returns in this context. Moreover, the practical application involves some difficulties. For example, the Ang and Bekaert (2002, 2004) explicitly “leave out many

aspects of international asset allocation that may be important”, such as transaction costs. Bauer et al. (2004) point out that substantial parts of the excess returns disappear after accounting for transaction costs, and Ammann and Verhofen (2006) find weak indications that their switching strategy remains profitable out-of-sample. The results of Hess (2006) veer toward the same direction: when taking transaction costs into account, CAPM strategies based on regime forecasts have no advantage w.r.t. a single-state benchmark. According to Hess (2006), two major reasons for the relatively poor performance of regime-switching models are (i) the inaccuracy of regime forecast, and (ii) noisy parameter estimates. This is in line with earlier findings of Michaud (1989) who considers mean-variance optimized portfolios to be “estimation error maximizers” in different context, and Dacco and Satchell (1999) who show that even a small number of wrong regime forecasts is sufficient to lose any advantage of a superior model. Besides, often only relatively short sequences of monthly data are available.

The alternative approach proposed in this study intends to circumvent the problems mentioned above and differs from the established techniques in three ways. Firstly, the analysis is based on daily instead of monthly data. On the one hand, this increases the amount of data available for markets with short history. On the other hand, the impact of wrong regime forecasts reduces from an entire month to a single trading day. As Hess (2006) stated, the “performance [...] crucially depends on the quality of the regime forecasts” and “a wrong regime forecast may not only lead to a non-optimal but to a detrimental allocation in the contrary direction relative to the ‘neutral’ single state. Several phases of outperformance relative to the standard formulation are necessary to make up the damage caused by one single wrong regime forecast”. As to the HMM selected, several authors employ the basic regime-switching model described in Equation (1) to model daily return series. The best-known article is presented by Rydén et al. (1998) who analyze the S&P 500 and, although various extensions of the original model exist, several studies still rely on the standard approach (e.g. Linne, 2002; Bialkowski, 2003). A common finding is that a two-state model with conditional Gaussian distributions often fits the data satisfactorily, while three-state-models exhibit the tendency to

allocate outliers in a separate regime. Preliminary analyses of our data confirm these findings, and therefore all models have two states.

Secondly, a mean-variance optimization strategy is not the focus of this study. The intention is to develop a highly robust approach, and to avoid difficulties related to the joint prediction of means and variances. According to Bauer et al. (2004), investors who use mean-variance optimization procedures face negative effects from the deviation of the estimated risk and return parameters from the true figures. Therefore, we concentrate on the second moment only. The strategy is straightforward, without loss of generality let $\sigma_1 < \sigma_2$:

1. For time t , estimate the hidden state \hat{s}_t
2. Determine the weights of the portfolio at time t . If $\hat{s}_t = 1$, invest 100% in the index X_t , else 100% in the risk-free asset (Cash).

The intention is to reduce overall portfolio risk during volatile market periods by shifting from equities into the risk-free asset class. Mainly two possibilities exist to estimate the hidden states. Either by global decoding using the Viterbi algorithm (Viterbi, 1967), a dynamic programming technique that calculates the most probable sequence of hidden states by

$$\{\hat{s}_1, \dots, \hat{s}_T\} = \underset{j_1, \dots, j_T}{\operatorname{argmax}} P(S_1 = j_1, \dots, S_T = j_T | X_1^T = x_1^T), \quad (2)$$


or alternatively by local decoding based on the smoothing probabilities

$$P(S_t = j | X_1^T) \quad \forall j \in \{1, \dots, J\}, t \in \{1, \dots, T\}.$$

In contrast to the Viterbi algorithm, the smoothing probabilities locally determine the probability of a sojourn in state j at time t . Thus, the resulting path derived from the maximum probabilities is the sequence of most probable states, which does not correspond with the most probable state sequence in general. A preliminary analysis and the in-sample results in Section 3.1 show that the estimated conditional variances are almost identical for either approach. However, the number of

Figure 1: Number of wrong state classifications by the Viterbi and smoothing algorithm

The figure shows the proportion of wrongly classified states per position in per cent. The position is identified by the observations number, ranging from 1 to 250.



regime-switches is significantly lower for the Viterbi paths, as this algorithm reacts less to fugacious regime changes, and thus transaction costs are reduced (see simulation study below). Therefore, the Viterbi algorithm serves for estimating the hidden states in what follows.

Thirdly, the out-of-sample forecasts are subject to a filtering procedure. This reduces undesired frequent state changes and thus transaction costs, Section 3.2 contains details on the filter. The following simulation study briefly illustrates the motivation, the underlying data are 50000 simulated series of length 250 each from a two-state HMM with parameters estimated from the S&P 500 data. For every series, estimated of the underlying state sequences are calculated by the Viterbi algorithm and the smoothing probabilities. For the latter, we assume that the underlying state has a low variance if $P(S_t = 1 | X_1^T) > 0.5$ and a high variance otherwise. Figure 1 displays the proportion of wrong state classifications per position. It is obvious that both techniques perform rather well for the larger part of the observations, with average errors of 3.48% and 3.16%. However, classification

Table 1: Descriptive statistics of daily returns

This table summarizes the daily returns data of the DAX, the DJIA, the NASDAQ 100, the Nikkei 225, and the S&P 500 index. It displays the number of observations, mean, standard deviation, skewness, kurtosis, and the value of the Jarque-Bera test statistic.

Name	N	Mean·10 ⁴	S.D.·10 ²	Skew.	Kurt.	JB
DAX	7,796	3.21	1.25	-0.426	10.26	17,384
DJIA	7,823	3.45	1.02	-2.144	59.09	1,032,042
NASDAQ 100	5,383	5.18	1.79	-0.089	10.04	11,142
Nikkei 225	5,925	1.42	1.31	0.034	7.56	5,136
S&P 500	7,834	3.54	1.00	-1.692	43.30	534,098

errors at the beginning and the end of a sequence increase strongly to circa 10% and 6.5%, respectively, independent of the chosen algorithm. This is problematic insofar as the last position plays a central role for the state prediction at time $T + 1$ in an out-of-sample setting. As mentioned already in the previous paragraph, an important feature from the practical point is that the Viterbi-path reduces the number of regime switches compared to smoothing (from 2.92 to 2.55). To further reduce the number of switches, a smoothing of the out-of-sample forecasts is presented in Section 3.2.

2 Data

The data analyzed in this paper are daily returns for five major equity indices, each covering at least 20 years: DAX, DJIA, NASDAQ 100, Nikkei 225, and S&P 500. The data for the DAX, DJIA, and S&P 500 start in January 1976, whereas the records of the NASDAQ and Nikkei begin in October 1985 and January 1983, respectively. Following Ang and Bekaert (2002), we fix the return of our risk-free asset to an annual rate of 3%. This figure can be considered being relatively conservative when compared to the previously cited study; it guarantees that the risk-free return is attainable during almost all periods and markets. Returns are calculated as $R_t = \ln(P_t) - \ln(P_{t-1})$, where P_t represents the index closing price on day t , adjusted for dividends and stock splits. Table 1 provides descriptive statistics for the data.

Table 2: Parameter estimates for the in-sample Markov-switching model with conditional Gaussian distributions

Estimated parameters for a model with Markov-switching Gaussian component distributions for the five indices. For state $i = 1, 2$, p_{ii} is the entry on the diagonal of the TPM, and μ_i and σ_i parameterize the conditional distributions, respectively.

	p_{11}	p_{22}	$\mu_1 \cdot 10^4$	$\mu_2 \cdot 10^4$	$\sigma_1 \cdot 10^2$	$\sigma_2 \cdot 10^2$
DAX	0.991	0.977	6.94	-6.62	0.786	2.01
DJIA	0.992	0.950	4.99	-6.66	0.776	1.98
NASDAQ 100	0.994	0.985	10.70	-9.22	1.090	2.90
Nikkei 225	0.980	0.989	10.60	-3.43	0.468	1.58
S&P 500	0.991	0.968	5.90	-5.31	0.712	1.67

All indices are leptokurtic and the Jarque-Bera statistic confirms the departure from normality for all series at the 1% significance level. Note that the year 1987 contains a unique event, the 'Black Monday' on October 19th 1987. On this day most of the indices sharply retreated, e.g., the DJIA and the S&P 500 lost 25.6% and -22.8%, respectively. In order to reduce manual interventions in our analysis, we omit an outlier correction. Consequently, the estimated values of the kurtosis, in particular of the two broad US indices, have to be looked at with caution.

3 Empirical Results

3.1 In-sample results

For every index, the Markov-switching model described in Section 1.1 is fitted. Table 2 summarizes the estimation results, which all display typical features common to return series: the two regimes are clearly separated with conditional variances differing by factor two to three and conditional means close to zero. Both regimes are highly persistent, whereas parameter estimates for the high-variance regime indicate that sojourns in this regime tend to last shorter than sojourns in the low-variance regime (with exception of the Nikkei).

Table 3 presents the in-sample performance of the asset allocation strategy described in the previous section. Additionally, it shows the performance of a strategy

Table 3: In-sample performance of Markov-switching strategies

This table displays annualized returns (in %), standard deviation (in %), and Sharpe ratios of the five indices and the Markov-switching strategies. “*Str.^{Vit.}*” and “*Str.^{Smo.}*”, respectively, denote the Viterbi- and smoothing-based strategy. Every index is followed by the statistics of the respective strategies in the two subsequent rows.


Name	Mean	S.D.	Sharpe ratio	Transitions
DAX	6.2	19.8	0.25	-
<i>Str.^{Vit.}</i>	13.8	10.7	0.98	60
<i>Str.^{Smo.}</i>	13.8	10.6	0.99	86
DJIA	7.6	16.1	0.35	-
<i>Str.^{Vit.}</i>	11.2	11.6	0.71	54
<i>Str.^{Smo.}</i>	11.4	11.5	0.74	74
NASDAQ 100	9.3	28.3	0.35	-
<i>Str.^{Vit.}</i>	19.8	14.8	1.10	28
<i>Str.^{Smo.}</i>	21.0	14.7	1.17	38
Nikkei 225	1.4	20.7	0.03	-
<i>Str.^{Vit.}</i>	11.6	4.3	1.87	46
<i>Str.^{Smo.}</i>	11.6	4.3	1.86	62
S&P 500	7.9	15.7	0.37	-
<i>Str.^{Vit.}</i>	12.4	10.1	0.91	58
<i>Str.^{Smo.}</i>	12.8	10.1	0.95	82

based on smoothing probabilities for comparability. The first three columns contain annualized mean, standard deviation, and Sharpe ratio, and the last column displays the number of transitions required by the strategies.

The main findings can be summarized as follows. Firstly, the exposure to high-volatile periods is strongly reduced. Secondly, the returns of the strategies are significantly higher and thus is the Sharpe-ratio. This side-effect results from periods of high volatility tending to coincide with periods of falling stock prices (Schwert, 1989). Thirdly, the performance of the two strategies does not differ significantly, apart from the number of transitions, as indicated in the previous section. Note that transaction costs are not taken into account in the results reported in Table 3. However, in the face of the observation period of several decades the number of transitions is almost negligible.

More importantly, it should be noted that the performance is strictly in-sample, and regarding Figure 2, the nature of these results becomes clear. All major declines

Figure 2: Indices with state sequence determined by the Viterbi algorithm. The five panels show plots of the daily returns from the five indices. The returns are shaded gray in the high-variance regime and black in the low-variance regime. The underlying state sequences result from Viterbi paths.



(e.g., stock market crash 1987, Asian crisis, Russian crises, crash of the dotcom bubble) coincide with periods of high volatility and, in particular, trading days ending with a strong decline regularly are linked to the high-variance state. Nevertheless, it confirms the result of Ang and Bekaert (2002) that the “cost of ignoring regime-switching is very high if the investor is allowed to switch to a cash position”.

3.2 Out-of-sample results

Aim of this section is to perform an out-of-sample forecast study under realistic conditions. In order to avoid transaction costs, the focus lies on the more steady Viterbi paths only, and transaction costs are fixed at 10 basis points (0.10%) for a one-way trade. While this might not be achievable for private investors, it represents a conservative assumption for professionals who can implement the proposed strategy in a very cost-efficient way using index future contracts.

The first step of the out-of-sample forecasts is implemented as follows. Select a window of n observations $[x_{t-n+1}, \dots, x_t]$, where x_t corresponds to the last available observation, and fit a HMM. Subsequently, calculate the Viterbi-path using the estimated parameters. Then, derive the probability of being in state $i \in \{1, 2\}$ at time $t + 1$ conditional on the knowledge of \hat{s}_t , which is a simple multiplication of the TPM with a vector of zeros except of a one at position \hat{s}_t (MacDonald and Zucchini, 2009). For the out-of-sample predictions, the length of the rolling window is set to $n = 2000$, equaling about eight years of historical data which can be thought of the average length of a full economic cycle.

In a second step, the sequence of predicted states is smoothed to further reduce the number of state shifts and thus transaction costs. Many alternative filtering procedures exist, but here we apply a simple median filter of lag k . That is, the predicted state at time $t + 1$ is given by

$$\hat{s}_{t+1}^f = [\text{median}(\hat{s}_{t-k+2}, \dots, \hat{s}_{t+1})], \quad (3)$$

Table 4: Effect of the median filter

The table summarizes the number of state changes before and after applying the median filter. The median filter includes the current and five past observations, i.e. $k = 6$. The original paths result from a Viterbi-based prediction with a window of length 2000.

Name	Before	After	# obs.
DAX	188	84	5,796
DJIA	190	60	5,823
NASDAQ 100	59	31	3,383
Nikkei 225	207	89	3,925
S&P 500	134	46	5,834

where $[\cdot]$ maps every number to its integer part. The states \hat{s}_{t+1-j} , $j = 0, \dots, k - 1$ correspond to the predicted state from the first step, each from a different window. Table 4 displays the effect of the filter with $k = 6$ on the number of transitions in the predicted state sequence. The number of state changes reduces by about 50-65%, which has a significant impact on the transactions costs.

For an application of the median filter, the value selected for the lag parameter k has non-negligible impact on the performance. A too high value delays the reaction to regime changes too much, whereas a too low value increases the number of transactions and therefore produces costs. The choice of six trading days produces good results for all indices, although further numerical optimization of k per index would improve the results. However, the primary goal is to develop a simple and robust strategy which simultaneously works for all markets, and therefore these techniques are not pursued here. Table 5 summarizes the out-of-sample performance of the strategies and corresponding indices taking transaction costs into account.

For all indices, the exposure to highly volatile periods is reduced. Investors following the strategy significantly reduce their risk in terms of the annualized standard deviation, on average by 41%. The highest degree of risk reduction is observed for the NASDAQ where the standard deviation of the strategy (14%) is not even half the risk of the index (32%). Moreover, all strategies outperform the respective index in terms of annual returns. The highest average annual excess return is realized for the Nikkei (201.6 bp), the lowest difference occurs for the S&P 500 (18.5 bp). This is naturally much lower than in-sample, however a not undesirable

Table 5: Out-of-sample performance of Markov-switching strategies

This table displays annualized mean returns (in %), standard deviations (in %) for the five indices and the Markov-switching strategies. Sharpe ratios are only reported for positive mean returns. “*Str.^{Vit.}*” denotes the Viterbi-based strategy. Every index is followed by the statistics of the respective strategy in the subsequent row.

Name	Mean	S.D.	Sharpe ratio	# Forecasts	# Transitions
DAX	7.24	22.0	0.292	5,796	-
<i>Str.^{Vit.}</i>	7.76	13.0	0.437	-	84
DJIA	8.93	16.8	0.417	5,823	-
<i>Str.^{Vit.}</i>	9.82	11.2	0.646	-	60
NASDAQ	6.82	32.0	0.272	3,383	-
<i>Str.^{Vit.}</i>	8.63	14.0	0.464	-	31
Nikkei	-4.30	22.6	-	3,925	-
<i>Str.^{Vit.}</i>	-2.28	13.9	-	-	89
S&P 500	8.37	16.5	0.390	5,834	-
<i>Str.^{Vit.}</i>	8.56	10.3	0.577	-	46

side-effect of avoiding volatile periods. Consequently, the strategies exhibit much better Sharpe ratios than the respective indices. These results are in line with Ang and Bekaert (2004) for monthly returns who noted that “optimal regime-switching asset allocation may require shifting assets into bonds or cash when a bear market is expected”. Although the regimes in this study do not explicitly try to identify bear markets by negative conditional returns, these periods are apparently circumvented by avoiding the high-variance regime.

4 Conclusion

The figures show that the answer to our initial question, whether profitable Markov regime-switching strategies exist, is a clear “Yes”. While previous works look at a range of different samples, they largely focus on monthly data, and often do not consider out-of-sample performance or transaction costs. The present paper contributes an investigation of the profitability of an asset allocation strategy that is based on a Markov-switching approach applied to daily returns. Looking at five major regional markets over the last four decades, the performance of the proposed model in a simple two-asset world is evaluated in- and out-of-sample under realistic

assumptions.

The in-sample analysis delivers two main results. At first, an asset allocation strategy, where the switching signals are derived either by the Viterbi algorithm or by smoothing, clearly outperforms a buy-and-hold strategy. While the returns of the Markov-switching based strategies are found to be higher, the corresponding volatilities are observed to be significantly lower than the volatilities of the respective indices. Secondly, the number of transitions in case where the prevailing market regime is derived by the Viterbi algorithm is lower than the number of necessary switches related to the smoothing-based strategy.

In an out-of-sample context with a focus on Viterbi-based strategies, the existence of a profitable Markov-switching based asset allocation strategy can be confirmed. Employing a robust technique to reduce the number of regime switches and thus transaction costs, the results are encouraging: For all analyzed stock market indices, the strategy is found to be profitable after transaction costs, which are taken explicitly into account. Portfolio risk is lowered by an average of 41% for all markets under consideration. Besides, annualized excess return between 18.5 (S&P 500) and 201.6 (Nikkei) basis points can be realized by avoiding highly volatile periods. This leads to Sharpe ratios between 0.437 (DAX) and 0.646 (DJIA), which compares well to an average Sharpe ratio of the various indices of 0.342. This is a significant improvement of former studies exerting monthly returns.

The methodology employed in this study can be extended in various directions. In a first step, the investable universe could be extended by also allowing other asset classes, such as bonds or commodities. It would be of interest to see how the observed profitability depends on the spectrum of available investment instruments. Another open route for future research relates to the use of more flexible models, such as semi-Markov models, which allow for nonparametric state occupancy distributions. By employing a regime-switching framework that is characterized by an even higher degree of flexibility, profitability of the proposed asset allocation strategy might be further improved. Consideration of economic covariates on the one hand and smoother switching procedures for the change from cash to shares and vice versa on the other hand may also be approaches that are worth further

exploration.

References

- Ammann, M. and M. Verhofen (2006). The effect of market regimes on style allocation. *Financ. Markets Portf. Manage.* 20(3), 309–337.
- Ang, A. and G. Bekaert (2002). International asset allocation with regime shifts. *Rev. Finan. Stud.* 15(4), 1137–1187.
- Ang, A. and G. Bekaert (2004). Timing and diversification: A state-dependent asset allocation approach. *Financial Analysts J.* 60(2), 86–99.
- Bauer, R., R. Haerden, and R. Molenaar (2004). Timing and diversification: A state-dependent asset allocation approach. *J. Investing* 13(3), 72–80.
- Baum, L. E. and T. Petrie (1966). Statistical inference for probabilistic functions of finite state Markov chains. *Ann. Math. Statist.* 37, 1554–1563.
- Baum, L. E., T. Petrie, G. Soules, and N. Weiss (1970). A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *Ann. Math. Statist.* 41, 164–171.
- Bialkowski, J. (2003). Modelling returns on stock indices for western and central european stock exchanges - Markov switching approach. *Southeast. Eur. J. Econ.* 2(2), 81–100.
- Brinson, G., B. D. Singer, and G. L. Beebower (1991). Determinants of portfolio performance ii: An update. *Fin. Anal. J.* 47(3), 40–48.
- Bulla, J. and A. Berzel (2008). Computational issues in parameter estimation for stationary hidden Markov models. *Computation. Stat.* 23(1), 1–18.
- Cappé, O., E. Moulines, and T. Ryden (2007). *Inference in Hidden Markov Models*. Springer Series in Statistics. New York - Heidelberg - Berlin: Springer-Verlag.

- Carhart, M. M. (1997). On persistence in mutual fund performance. *J. Finance* 52(1), 57–82.
- Dacco, R. and S. Satchell (1999). Why do regime-switching models forecast so badly? *J. Forecasting* 18(1), 1–16.
- Dempster, A. P., N. M. Laird, and D. B. Rubin (1977). Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Statist. Soc. Ser. B* 39(1), 1–38. With discussion.
- Goldfeld, S. M. and R. E. Quandt (1973). A markov model for switching regressions. *J. Econometrics* 1(1), 3–16.
- Graffund, A. and B. Nilsson (2003). Dynamic portfolio selection: The relevance of switching regimes and investment horizon. *Europ. Finan. Manage.* 9(2), 179–200.
- Guidolin, M. and A. Timmermann (2005). Economic implications of bull and bear regimes in UK stock and bond returns. *Econ. J.* 115(500), 111–143.
- Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica* 57(2), 357–384.
- Hess, M. K. (2006). Timing and diversification: A state-dependent asset allocation approach. *Europ. J. Finance* 12(3), 189–204.
- Linne, T. (2002). A Markov switching model of stock returns: an application to the emerging markets in central and eastern europe. In *in: East European Transition and EU Enlargement*, pp. 371–384. Physica-Verlag.
- MacDonald, I. L. and W. Zucchini (2009). *Hidden Markov for Time Series: An Introduction Using R*. CRC Monographs on Statistics and Applied Probability. London: Chapman & Hall.
- Michaud, R. O. (1989). The markowitz optimization enigma: Is 'optimized' optimal? *Financial Analysts J.* 45(1), 31–42.

- Quandt, R. E. (1958). The estimation of the parameters of a linear regression system obeying two separate regimes. *J. Amer. Statistical Assoc.* 53(284), 873–880.
- Quandt, R. E. (1972). A new approach to estimating switching regressions. *J. Amer. Statistical Assoc.* 67(338), 306–310.
- Rabiner, L. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *IEEE Trans. Inf. Theory* 77(2), 257–284.
- Rydén, T., T. Terasvirta, and S. Asbrink (1998). Stylized facts of daily return series and the hidden Markov model. *J. Appl. Econom.* 13(3), 217–244.
- Schwert, G. W. (1989). Why does stock market volatility change over time. *J. Financ.* 44(5), 1115–1153.
- Turner, C. M., R. Startz, and C. R. Nelson (1989). A Markov model of heteroskedasticity, risk, and learning in the stock market. *J. Finan. Econ.* 25(1), 3–22.
- Viterbi, A. J. (1967). Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Trans. Inform. Theory* 13(2), 260–269.