ИКОНОМИЧЕСКИ ИЗСЛЕДВАНИЯ
ECONOMIC STUDIES

Volume 32(1), 2023

Anton Gerunov[1]

STOCK RETURNS UNDER DIFFERENT MARKET REGIMES: AN APPLICATION OF MARKOV SWITCHING MODELS TO 24 EUROPEAN INDICES[2]

This paper studies the different modes of operation of European stock markets. Using data on 24 European indices over a period of 15 years, we show that these can be well represented by a Hidden Markov Model with two regimes that roughly correspond to bull and bear markets. We further estimate regime parameters and show that the alternate regimes have very different risk-return tradeoffs with clear implications for portfolio management. Corresponding transition probability matrices show the remarkable persistence of states and give a possible quantitative estimate of the degree of inertia in financial markets. Regime-switching coordination across markets is further examined, showing that moments of correlations are followed by idiosyncratic episodes and thus, risk diversification through regime arbitrage is possible.
Keywords: market returns; Markov switching model; regime change; European stocks
JEL: G11

1. Introduction

Estimating key financial market parameters is crucial in both theoretical and practical terms. On the academic front knowing the average returns, risks, and the risk premium allows the researcher to estimate risk aversion, the cost of capital and to model wider macroeconomic dynamics. For the practitioner, those are important when making an investment, savings, and retirement decisions. The usual approaches to reach those estimates focus on detailed time series of either realized returns or calculated discounted cash flows, with a spruce of more exotic approaches in-between. While all those methods have their benefits, it is very often that they assume a certain consistency of the estimated parameters. However, it is well established that financial markets fail to remain time-invariant and are subjects to episodes of expansions and contractions, disparate modes of volatility, and episodes of irrational (or rational) exuberance. In short, there are different regimes under which stock markets tend to operate (Baltas, Karyampas, 2018).

[1] *Anton Gerunov, Ph.D., D.Sc., Associate Professor, Sofia University "St. Kliment Ohridski", e-mail: a.gerunov@feb.uni-sofia.bg .*

[2] This paper should be cited as: *Gerunov, A. (2023). Stock Returns Under Different Market Regimes: An Application of Markov Switching Models to 24 European Indices. – Economic Studies (Ikonomicheski Izsledvania), 32(1), pp. 18-35.*

This article aims to investigate further whether this is the case for European stock markets by applying a data-driven approach to estimating different market regimes. To this end, we leverage a Hidden Markov Switching model that enables the calculation of different regimes, the density functions of their respective key parameters, and the switching probability throughout each. This is then used to estimate the long-run returns, risks, and equity risk premia in 24 European stock markets over a 15-year period ranging from 2006 through 2020. The results outline the significant difference between the financial markets with crucial implications for equity investors. Moreover, we show that highly developed financial markets tend to cluster together, while less developed ones display more or less idiosyncratic dynamics. This finding leads to distinct implications for practitioners and opens new venues of research.

The article is structured as follows. The second section begins by outlining how the return, risk, and equity risk premium (ERP) are correlated. The most common approaches to calculating the ERP are outlined, together with more recent results. The third section introduces the methodology and the model estimation approach for calculating different regimes by briefly outlining the Markov Switching Model and its relevance for financial research. The fourth section proceeds to make an overview of the 24 European financial markets under study by summarizing a few key stylized facts about them. The following section calculates Markov switching models for each of the markets and presents the numeric estimates for the realized returns, levels of risk, and equity risk premia. Then we elaborate on the results by outlining a few recommendations. The final section is a concluding one.

2. Essence and Application of Hidden Markov Models to Financial Markets

While it is widely acknowledged that stock market characteristics such as risk and return vary over time, with a clear tendency for similar observations to cluster together. However, the precise mechanics of regime switching are yet to be fully resolved. The fact that stock markets operate under different regimes is hardly novel – at a minimum their dynamics differ significantly in times of growth and in times of recession. Classical ideas, such as the Efficient Market Hypothesis (EMH), would imply random price movement spurred by new information. It remains unclear whether this information signals a switch to a new state of operation (regime) of the market or is just regular innovation within the current regime. This clearly puts forward the need to leverage a data-driven approach to decide what regime is currently the stock market operating in and what are the crucial characteristics of this specific state of nature.

To this end, the researcher may use a Hidden Markov Model, HMM (Nguyen, Nguyen, 2015). The idea behind such a model is that we have an objective state (e.g. jobless recovery, irrational growth, fear of recession) which is not directly visible in the data, but it leads to the generation of observable data points such as asset returns or levels of volatility. The Hidden Markov Model uses observable data to estimate the current regime and a matrix of transition probabilities for regime switching. In a similar vein, Nguyen (2017) proposes using an HMM for the prediction of technology stocks and finds that is effects significantly both stocks trading and derivatives hedging. While the HMM approach is well-recognized among other more familiar ones (Somani et al., 2014; Sutkatti, Torse, 2019), it still remains a relatively new addition to the financial toolbox as compared to more traditional analytic methods.

As early as 1998, Ryden et al. (1998) leveraged an HMM approach to model daily returns from 10 subseries of the S&P500 index. They use a combination of information criteria and reach the conclusion that HMMs estimated on their data tend to be characterized by either two or three regimes. Moving beyond those numbers was considered computationally prohibitive at the time. Those estimated HMMs are found to be a very good fit for S&P500 dynamics. In 2005 Hassan & Nath (2005) already noted that the HMM might be a new paradigm for analyzing financial markets. They model the price of interrelated airline stock using a Hidden Markov Model and utilize it for price forecasting. The main result is that this approach presents a viable forecasting mechanism.

Much research followed with Hidden Markov Models being applied to a wide variety of pertinent theoretical, as well as practical issues. Hassan (2009) further expands on HMM modelling by proposing to combine the HMM with a fuzzy model for stock market forecasting and apply it to data on six stocks (three airlines and three IT companies). The combined model is shown to outperform competing ones such as an Autoregressive Moving Average (ARIMA), Artificial Neural Network (ANN), and another HMM. Similarly, Gupta and Dhingra (2012) study 4 stocks and found that an ensemble HMM outperforms both ARIMA and ANN. As an input parameter to the HMM they assume that the number of states is four, but this lacks formal justification. Zhang et al. further enhance the model by proposing their Extended Coupled Hidden Markov Model that is able to take into account news events as well as historical trading data and stock correlations. Testing it on Chinese A-share market data in 2016 shows that it outperforms a number of state-of-the-art alternatives.

Nguyen and Nguyen (2015) use approach the portfolio selection problem with a sophisticated HMM-based approach. They make monthly regime predictions for the consumer price index, industrial production, S&P500, and market volatility. In parallel, they analyze all the stocks in the S&P500 index and assign them scores depending on how much they benefit from different regimes. Using this algorithm, they compose a portfolio of stocks which turns to significantly outperform the benchmark S&P500 – by generating a return of 14.9% to 2.3% for the index.

Baltas and Karyampas (2018), find strong effects on the specific regime of operation of the stock market on the valuation of the equity risk premium. They further argue that taking into account the regime dependence of returns allows one to better forecast the equity risk premium, thus improving the portfolio allocation decisions. In a 2018 paper, Nguyen (2018) applies a Hidden Markov Model to data on the closing prices of the US index S&P 500. Using a set of four information criteria, he finds out that an HMM with four regimes best captures the data at hand. This model is able to outperform traditional stock price prediction methods, thus underlining the viability of using a Hidden Markov Model for modelling and forecasting the stock market. In a similar vein, Kole (2019) uses S&P500 data to demonstrate the application and utility of an HMM model. Here, a 2-state HMM is assumed with little recourse to the need of precisely choosing the number of regimes of the market. In general, academics and practitioners tend to prefer fewer number of regimes, making for a more tractable model. Thus, a usual choice is to assume just two states (bearish/bullish market). It is the rare exception that a large number of regimes are assumed as in Sasikumar & Abdullah (2016).

Kim et al. (2019) use a Hidden Markov Model to identify the regime for global assets divided into 10 classes, or 22 subclasses, leveraging data from the beginning of 2004 to the end of 2018. The number of regimes is assumed to be three, with little recourse as to why this precise number is taken. They show that using the HMM improves portfolio results in a number of dimensions and propose that HMM may enable more stable portfolio management compared to momentum strategies. Acula and Guzman (2020) use an enhanced HMM model to predict the closing prices of Nokia and Apple stocks. In the former case, their proposed model outperforms a neural network, while in the latter one, they tend to be on par.

Dias et al. (2015) are among the few research papers that focus on European stock market dynamics over the period 1998-2013. They analyze data on 21 European stock market indices leveraging an extended Mixture Gaussian Hidden Markov Model, where they use the Bayesian Information Criterion (BIC) to select the number of model regimes. Their main conclusion is that in Europe, there are two groups of economies, and the optimum number of regimes across them is three. The first group tends to consist mainly of countries with developed and well-functioning financial markets, while the second consists of ones with less developed or recently negatively affected markets. The question of whether all countries are appropriately modelled by a three-state HMM is left to further research.

There are also some applications that leverage stock market data to detect systemic financial market phenomena. De Angelis and Paas (2013) use an HMM to detect financial crises and mark their ends in order to support investors' decision-making. Their proposed HMM is found to outperform an alternative state-of-the-art threshold General Autoregressive Conditional Heteroscedasticity (GARCH) model, thus showing the superiority of the HMM approach. Shi & Song (2012) use propose using an Infinite Hidden Markov Model (iHMM) to detect speculative bubbles. They show the utility of this approach by applying it to the Argentinian money aggregates over the period 1983 to 1989, as well as to the US oil prices from 1983 through 2010. Another line of applications of HMM includes the detection of stock price manipulation (Cao et al., 2013; Cao et al., 2014).

3. Multiple Regimes and Markov Switching

3.1. General Overview of the Hidden Markov Model

This paper leverages a relatively long data series spanning over 15 years of daily market data in 24 European markets to estimate a Hidden Markov Model and thus be able to discern the different states in which those markets have functioned. Here we present a short overview of the HMM, based on Nguyen (2018). The Hidden Markov Model assumes a number of states N and observed data of length T, which we denote as O:

$$O = \{O_t, t = 1, 2, 3, \dots, T\} \tag{1}$$

There is also a sequence of hidden states (Q) and a number of possible values for each state, S, where those are defined as follows:

$$Q = \{q_t, t = 1, 2, 3, \dots, T\} \text{ and } S = \{S_i, i = 1, 2, 3, \dots, N\} \quad (2)$$

There are v_k symbols per given state that characterize it. The model calculates the likelihood of realized observations and stores this in the observation probability matrix **B**, where:

$$\mathbf{B} = (b_{ik}): b_{ik} = P(O_t = v_k | q_t = S), i = 1 \dots N, k = 1 \dots M \quad (3)$$

As a final step, the model characterizes the probabilities of staying in the current state or switching to another one. Those are summarized in the probability transmission matrix, **A**. This is defined as follows:

$$\mathbf{A} = (a_{ij}): a_{ij} = P(q_t = S_j | q_{t-1} = S_i), i = 1 \dots N, j = 1 \dots N \quad (4)$$

The hidden Markov model can be thus characterized by the probability transmission matrix, the observation probability matrix and the vector of initial probabilities of being in a given state (denoted p). A compact notation of the Hidden Markov model λ is as follows:

$$\lambda \equiv \{A, B, p\} \quad (5)$$

With a continuous distribution of probabilities, as is the case with stock market prices and returns, the resulting Hidden Markov Model is also a continuous one. Should the observation probabilities data be normally distribution, then we have the following:

$$b_i(O_t) = N(O_t = v_k, \mu_i, \sigma_i) \quad (6)$$

Denoting the mean and standard deviation of this distribution as μ and σ, respectively, the Hidden Markov model in Eq. (5) reduces to:

$$\lambda \equiv \{A, \mu, \sigma, p\} \quad (7)$$

There are three key HMM estimation problems to be solved. The first one is calculating the probability of observations $P(O|\lambda)$ given observed data and model parameters. This is usually tackled by forward or backward algorithms (Baum, Egon, 1967; Baum, Sell, 1968). The second problem is to calculate the best state sequence of observations $Q = \{q_1, q_2, \dots, q_T\}$ again given the data at hand and the model parameters. This is accomplished by the Viterbi algorithm (Viterbi, 1967). The third and final calculation issue for HMMs is to calibrate the model parameters $\lambda = \{A, B, p\}$ using the available data. This is accomplished by applying the Baum-Welch algorithm (see Baum, Petrie, 1966). The latter is then extended in work by Levinson et al. (1983) and Li et al. (2000). Finding numeric solutions to those three problems can sometimes be challenging. To this end, we use the Expectation Maximization (EM) algorithm, proposed by Hamilton (1989, 1990). Essentially, the EM method maximizes the likelihood function of the HMM that can be used in general cases (Mizrach, Watkins, 1999). This method was later expanded and elaborated by Goldfeld & Quandt (2005) and developed for use in the R language for statistical programming by Sanchez-Espirages & Lopez-Moreno (2015). It is precisely this implementation that we use to estimate our HMM parameters.

3.2. Selecting the Number of Regimes

Hidden Markov Models can be a powerful tool to develop a subtler and more nuanced view of how markets operate by understanding their current state (or regime), estimating their parameters, and then reacting accordingly. While this can be a purely data-driven process, the key question of how many regimes there are remaining open. Preferably, the researcher will be able to estimate that from data rather than simply assume a plausible number. Still, some work in the field has taken the second route and posited that a likely number of states might very well be two, accounting for growth and fall conditions. While informative, one may want to further study whether a higher-order HMM is a better fit for data.

Noting the importance of the likelihood function for this end, some authors have proposed to use of a visual approach to understanding this. Plotting the likelihood function for a number of regimes, an observable kink in the line signals the desired order. This is somewhat informal and can be rather imprecise, which is why the usual way to approach regime selection is by resorting to information criteria (Pohle et al., 2017). Essentially, an information criterion is a measure of the quality of a given model based upon the informational loss this model entails. It is best used to compare alternative models fit on the same dataset and allows the researcher to take heed of the tradeoff between model quality (fit) and parsimony (number of model parameters). As Hidden Markov Models with a different number of regimes are calculated, one can estimate their criteria and select the ones with the lowest values.

Among the most commonly used criteria are the Akaike Information Criterion (AIC), the Consistent Akaike Information Criterion (CAIC), the Bayesian Information Criterion (BIC), and the Hannan-Quinn Information Criterion (HQC). Either one or more of those can be used to determine the number of states and this may also need to take into account the specific model context, the need for parsimony and issues of computational ease (Ngyuen, 2018). The likelihood function of the model shows the probability of observing the data given this specific model with its parameter set. This is denoted as L. We denote the number of observations with M, and the number of parameters in the parameter set – with k. If the distribution corresponding to each of N hidden states is a normal one, then the following holds true:

$$\lambda \equiv \{A, \mu, \sigma, p\} \tag{8}$$

The Akaike criterion (AIC) is then defined as follows:

$$AIC = -2\ln(L) + 2k \tag{9}$$

AIC has some issues, most notably that it is valid asymptotically and, in some situations, adjustments need to be made to take into account finite sample sizes. An alternative to AIC is the consistent AIC (CAIC) which is then defined as follows:

$$CAIC = -2\ln(L) + k(\ln(M) + 1) \tag{10}$$

Another popular and well-used information criterion is the Bayesian one (BIC). It should be noted that BIC is valid as M≫k, which is clearly the case with the current data under study. The BIC definition is given by the following:

$$BIC = -2\ln(L) + k\ln(M) \quad (11)$$

Finally, the Hannan-Quin criterion may be used for model selection. While it may not be asymptotically efficient, it is still well-behaved and consistent – and thus, a viable choice. The HQC is given by the following:

$$HQC = -2\ln(L) + k(\ln(M)) \quad (12)$$

In line with the literature and in order to ensure consistency and comparability of results, this paper takes a two-pronged approach. Initially, we assume a simple 2-state HMM and estimate parameters (returns and risk) for the European markets within this framework. Second, we fit a number of Markov Models and choose the optimal ones given the values of the Akaike and Bayesian information criteria (see Eqs. (10) and (11), respectively). The next section proceeds with a preview of data used for this fit and outlines a few key facts about the stock markets under investigation.

4. Data and Stylized Facts for European Stock Markets

European stock markets have had tumultuous dynamics over the period from 2006 to 2020. We leverage publicly available data from the Stooq platform spanning a 15-year period beginning 1st January 2006 and ending 31st December 2020 to study them in more detail. The focus is on 24 selected European stock market indices (see Table 1), each of them having full data for the period – a total of 3838 observations each. Data is automatically downloaded and processed via replicable scripts in the R programming language for both visualization and model estimation.

It should be noted that changing the span of the time series will likely generate slightly different parameter estimates. However, the period 2006-2020 seems relatively representative as it covers both periods of growth, as well as two large crises – the global financial crisis of 2007-2008 and the COVID pandemic that began in early 2020. The sample under study thus covers at least one full economic cycle.

In contrast to US-based indices, which showed an overall upward trend despite the global financial crisis that started in 2007-8, and the coronavirus pandemic of 2020, some European indices never recovered their levels from the beginning of the period (see Figure 1). Developed and sophisticated European economies such as Germany, France, and the United Kingdom have broadly followed the global stock market trends – after an initial fall in 2008, they mostly recouped the losses and registered growth as the economic recovery was gaining traction. The coronavirus pandemic in early 2020 led to a steep decline, but it has been mostly offset by gains in the next months. This is hardly the case with embattled economies in the South.

Figure 1. Dynamics of Selected European Stock Markets, 2006-2020

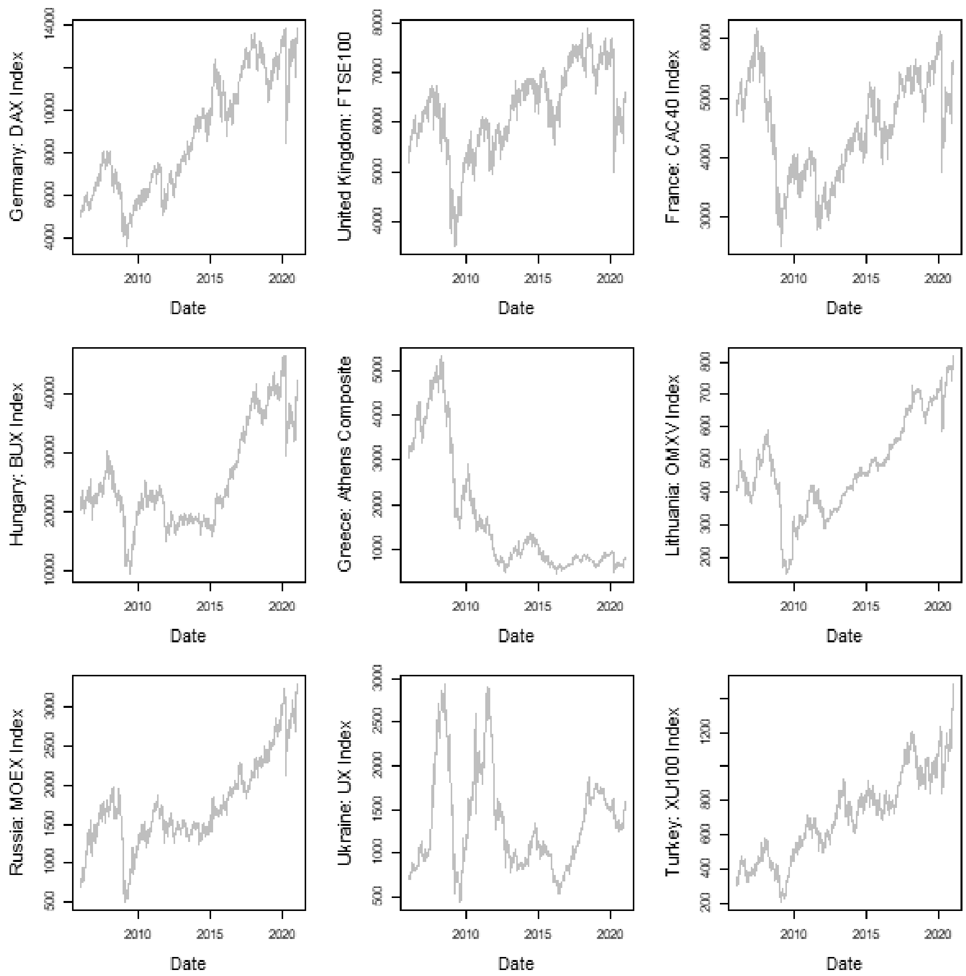

Source: Stooq Database.

Countries like Greece and Portugal experienced steep declines that were never recouped, and their stock markets have largely seen stagnation. Finally, a third group of countries that lies at the periphery of the European Economic Area (EEA) is characterized by yet another type of dynamics. Countries such as Russia and Turkey have registered a robust growth over the past fifteen years and have been largely unphased by the occasional declines driven by a deterioration in global fundamentals. Finally, Ukraine has been plagued by a constellation of heightened political, military and economic risks and has specific individual dynamics – twin peaks of growth followed by an unsteady recovery.

Table 1. Risk, return and simple equity risk premia in selected European stock markets

Country	Exchange	Index	Average Return, %	Equity Risk Premium, %	Risk (Std. Dev.), %
Belgium	Euronext Brussels	BEL20	2.91	1.42	21.93
Bulgaria	Bulgarian Stock Exchange	SOFIX	3.57	2.07	32.35
Czechia	Prague Stock Exchange	PX	-0.03	-1.52	19.91
Estonia	Estonian Stock Exchange	OMXT	10.22	8.73	31.95
Finland	Helsinki Stock Exchange	HEX	4.65	3.15	21.08
France	Euronext Paris	CAC	2.79	1.29	17.76
Germany	Deustche Boerse	DAX	8.40	6.91	19.61
Greece	Athens Exchange, Athens Composite	ATH	-2.84	-4.33	34.25
Hungary	Budapest Stock Exchange	BUX	8.88	7.39	29.31
Italy	Borsa Italiano	FMIB	-0.80	-2.29	20.38
Latvia	Latvian Stock Exchange, Riga	OMXR	6.83	5.33	24.94
Lithuania	Vilnius Stock Exchange	OMXV	8.88	7.39	28.07
Netherlands	Amsterdam Stock Exchange	AEX	4.73	3.24	19.77
Norway	Oslo Stock Exchange	OSEAX	9.97	8.47	23.18
Poland	Warsaw Stock Exchange	WIG20	0.51	-0.98	21.49
Portugal	Euronext Lisbon	PSI20	-0.77	-2.27	23.19
Romania	Bucharest Stock Exchange	BET	8.13	6.63	28.97
Russia	Moscow Exchange	MOEX	15.70	14.21	41.03
Spain	Bolsa de Madrid	IBEX	-0.06	-1.56	19.26
Sweden	Stockholm Stock Exchange	OMXS	6.43	4.94	19.62
Switzerland	Swiss Exchange	SMI	3.55	2.06	15.65
Turkey	Istanbul Stock Exchange	XU100	16.36	13.84	38.54
Ukraine	Ukraine Stock Exchange	UX	18.09	16.59	55.93
United Kingdom	London Stock Exchange, FTSE 100	UKX	1.91	0.42	13.90

Source: Stooq Database and author's calculations.

Fundamental market characteristics also vary widely across European stock markets (see Table 1). The average annual returns vary from the highs of Ukraine (18%), Turkey (16%), and Russia (15.7%) to negative average returns over the period as a whole – here, Greece particularly stands out with a -2.8% but a number of other stock market indices in other countries also register long-run returns below zero – Italy (-0.8%), Portugal (-0.8%), Slovakia (-0.3%) and Spain with a return around zero. The risk-return tradeoff holds quite nicely in European data, with the highest levels of return also being accompanied by higher standard deviations – a proxy for risk. The largest standard deviations are expected to be registered for Russia, Turkey, and Ukraine. However, almost all European markets are characterized by high levels of volatility which are sometimes set against the backdrop of relatively modest positive or even negative annual returns.

Table 1 also presents an indicative estimate of the equity risk premium across the European market. We calculate it as the difference between the returns of a reference riskless asset and the actual index returns (Damodaran, 2020). While this approach has a number of well-known potential complications (see Siegel, 2017 and references therein), it still serves as a reasonable proxy for the actual equity risk premium. A potential issue here is the determination of the riskless return. Considering that European stock markets are quite open (especially within the European Union) and almost all investors do have access to buying German bonds, we select the German government's 10-year treasury bond yield as the reference. We use the daily quotes of the bond, thus allowing for the daily dynamics of the risk-free rate proxy over the period under study.

The calculated equity risk premia imply that there are drastically different levels of risk across different European stock markets. This is in consonance with what one concludes from surveying the volatility levels of the indices. While it is clear that those results are partly driven by the global financial crisis of 2007-2008, the ensuing sluggish recovery, and the coronavirus pandemic that hit Europe hard at the beginning of 2020, they are markedly different from the dynamics of other developed economies. This merits further study and necessitates the development of a more nuanced understanding of how stock markets in Europe function. A final important fact for European stock markets is that they are highly interconnected. Investigating the correlational structure between annual index returns reveals high relations across practically all markets in Europe (see Figure 2). Those correlations are medium to large in size and statistically significant. The only notable exception is the Istanbul stock market index which shows negative correlations with most other stock markets. The strong links between markets indicate some success in Europe's economic integration goals. On the other hand, risk hedging is made more difficult as all markets tend to move in the same direction, thus diminishing the benefits of diversification across different geographies.

Figure 2. Annual return correlations across selected European stock markets, 2006-2020, grouped via hierarchical clustering

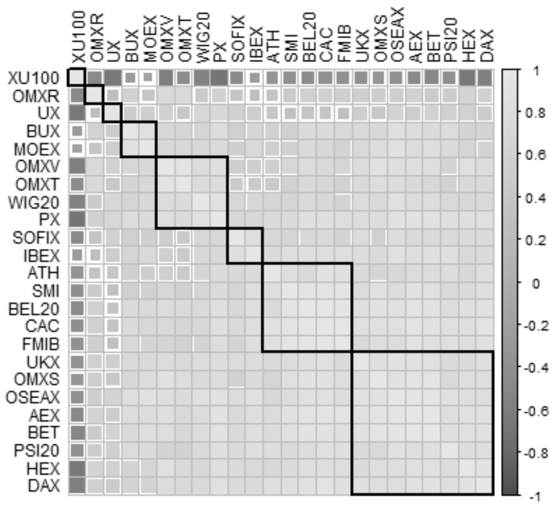

Source: Stooq Database and author's calculations.

In order to outline clusters of highly correlated markets, we grouped index correlations using hierarchical clustering. This exercise reveals three larger clusters of stock markets that tend to be intimately connected. The first one corresponds to central European countries and encompasses the market indices in Poland, the Czech Republic, Estonia, and Latvia. The second cluster approximately reflects the tight relationship in the centre and South of the continent around the gravity of the French and Italian economies and is underlined by the high correlations between the stock exchanges in France, Greece, Italy, Switzerland and

Belgium. The final cluster contains the most developed economies in western Europe and is dominated by Germany, the Netherlands, the United Kingdom, and the Nordics.

Some more isolated markets, such as the stock exchanges in Turkey, Russia, Ukraine, and a few others, do not belong to any clearly defined cluster. Those insights roughly correspond to the results obtained by Dias et al. (2015), whereby they find three large groupings of economies on the European continent, also distinguished by their levels of economic development and financial markets sophistication. However, one needs to also keep in mind that the correlations between markets hold over a longer period but not necessarily over a shorter one. Figure 2 presents the correlational structure of daily returns and here, we observe much less linkage and dependence as contrasted to the correlational structure of annual returns. The large cluster of central and western European markets persists, as does the cluster of southern markets. The rest turn out to be much less linked.

Figure 3. Daily return correlations across selected European stock markets, 2006-2020, grouped via hierarchical clustering

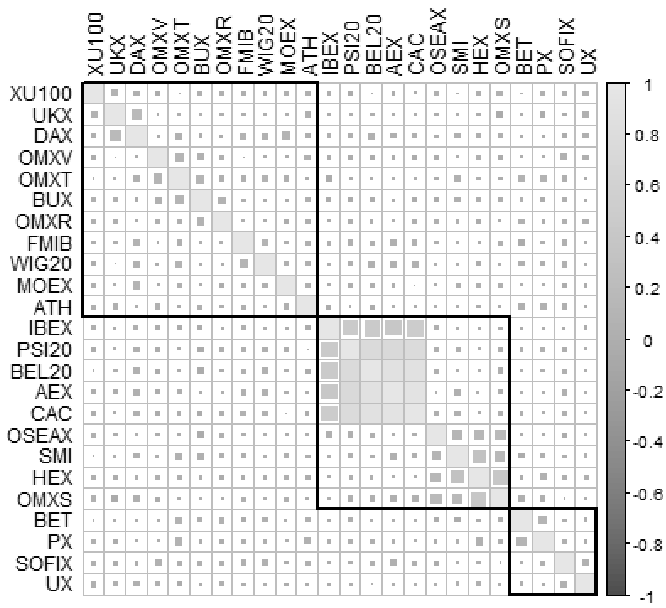

Source: Stooq Database and author's calculations.

Those stylized facts spell a very nuanced risk landscape in Europe and call for a more careful approach to risk management and portfolio choice. High correlations of annual returns show that investors can only achieve diversification and risk reduction in European markets if they make portfolio optimization decisions at a much lower level of granularity – e.g. on a daily basis where the market is not so tightly synchronized. The time structure of these correlations may need to be exploited as some markets react with a lag to overall downturns. All this points clearly to the need to understand daily market dynamics and be able to discern the current regimes they operate in, and the probability of switching to another one. The next section presents the application of a Hidden Markov Switching Model to these ends.

5. Estimating Regime Number and Parameters

European stock markets have had tumultuous dynamics Data on 24 European stock market indices is leveraged to fit the optimal Markov Switching model on those. Following the guidelines for parsimony, we calculate 7 different models for each index with a number of different regimes varying from 2 to 8. Those models are then compared against each other using the Akaike and Bayesian information criteria. With almost no exception, both criteria give favour to models with a very limited number of regimes showing the lowest values for 2-regime models. This makes both intuitive sense as markets are colloquially described as bear and bull markets, implying two modes of operation, and is in consonance with a large body of research literature. Since the two regimes calculated by the HMM may not completely cover the definitions of bull and bear markets but merely approximate their qualitative properties, we may instead call them bullish and bearish market regimes. Thus, the bullish market regime is expected to generate higher returns with lower volatility with respect to its bearish counterpart. The second analytic step is to estimate a 2-regime Hidden Markov Model for each market, and their respective parameters and switching probabilities are calculated using consequently the forward algorithm for computing the probability of observations, the Viterbi algorithm for estimating the best state sequence, and the Baum-Welch algorithm for calibrating model parameters.

Figure 4. Regimes in the FTSE100 index returns: bullish and bearish periods (top panel) and regime probabilities (bottom panel)

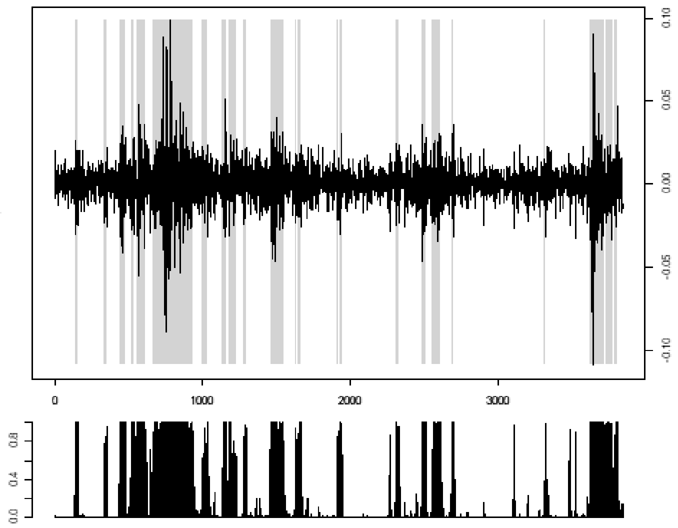

Source: Stooq Database and author's calculations

The HMM is able to designate each daily return under study as belonging to one of the two regimes with a different probability. An example case in point would be the UK's FTSE100 index, visualized in Figure 3. Most of the time, the market is in Regime 1 (bullish) with positive average returns of 0.05% daily. However, there are some turbulent times during which the market switches to the bearish Regime 2 (shaded regions in Fig. 4) where average daily returns turn negative to a -0.10% and the volatility increases significantly by two and a half times. The mode results are derived with high fidelity as most regime switches are based on high probabilities and not on a decision around the cut-off point. Similar dynamics can be observed across all European stock markets (see Table 2).

Under the bullish market regime, daily returns in European stocks are positive and mostly fall in the range of 0.05% to 0.1%, with relatively low volatility as proxied by standard deviations of 0.4% to around 1.2%. However, as the regime changes, market dynamics become very different. The average returns notably decline and, in all but one market, turn negative. Surprisingly, the lowest falls are not necessarily observed only in the riskiest markets – the steepest decline is seen in the Istanbul stock exchange, with average daily returns reaching -0.24%, followed by the Greek stock exchange (-0.20%), the Prague stock exchange (-0.19%) and the Swiss stock exchange (-0.17%).

Table 2. Risk and return under two different regimes in selected European stock markets

Exchange	Index	Regime 1: Bullish Market (%)		Regime 2: Bearish Market (%)	
		Returns	*Std. Dev.*	*Returns*	*Std. Dev.*
Euronext Brussels	BEL20	0.07	0.77	-0.14	2.06
Bulgarian Stock Exchange	SOFIX	0.02	0.58	-0.11	2.17
Prague Stock Exchange	PX	0.05	0.81	-0.19	2.64
Estonian Stock Exchange	OMXT	0.04	0.49	-0.02	1.85
Helsinki Stock Exchange	HEX	0.08	0.86	-0.14	2.14
Euronext Paris	CAC	0.07	0.86	-0.12	2.20
Deustche Boerse	DAX	0.09	0.92	-0.13	2.29
Athens Exchange, Athens Composite Index	ATH	0.07	1.25	-0.20	3.09
Budapest Stock Exchange	BUX	0.06	0.99	-0.07	2.51
Borsa Italiano	FMIB	0.05	1.10	-0.16	2.70
Latvian Stock Exchange, Riga All-shares Index	OMXR	0.04	0.70	0.01	2.38
Vilnius Stock Exchange	OMXV	0.05	0.42	-0.07	1.90
Amsterdam Stock Exchange	AEX	0.07	0.81	-0.14	2.21
Oslo Stock Exchange	OSEAX	0.09	0.90	-0.12	2.43
Warsaw Stock Exhange, 20	WIG20	0.04	1.05	-0.08	2.36
Euronext Lisbon	PSI20	0.06	0.84	-0.21	2.06
Bucharest Stock Exchange	BET	0.06	0.72	-0.06	2.56
Moscow Exchange	MOEX	0.09	1.08	-0.09	3.72
Bolsa de Madrid	IBEX	0.05	0.96	-0.11	2.35
Stockholm Stock Exchange	OMXS	0.08	0.84	-0.09	2.15
Swiss Exchange	SMI	0.07	0.74	-0.17	2.00
Istanbul Stock Exchange	XU100	0.12	1.21	-0.24	2.72
Ukraine Stock Exchange	UX	0.07	1.05	-0.06	3.46
London Stock Exchange, FTSE 100	UKX	0.05	0.72	-0.10	1.94

Source: Author's calculations.

The great difference between the states is also manifest itself in the very different volatility values. Under the bullish market, we observe volatilities of around or under 1%, with the highest being in Athens (1.25%), Istanbul (1.21%) and Italy (1.10%). However, as the market turns bearish, the volatility increases three or four times and its average value jumps from 0.86% into the calm Regime 1 all the way to an average of 2.41% under Regime 2. Embattled countries that have witnessed economic problems over the period of study consistently register the highest stock market index volatilities under the bearish regime. The highest standard deviation is registered in the Russian stock market (3.76%), followed by the Ukrainian and Greek ones with 3.46% and 3.09%, respectively. Remarkably, across Europe, there is really no safe haven during bearish markets – all the rest of the indices register high fluctuations in the range of 2-3% daily under Regime 2. This shows that while bearish markets may vary widely across European stock exchanges, bullish markets have strikingly similar characteristics in terms of volatility.

Table 3. Transition probabilities in two different regimes in selected European stock markets

Exchange	Index	Regime 1: Bullish Market (%)		Regime 2: Bearish Market (%)	
		Prob. to Stay	*Prob. to Switch*	*Prob. to Stay*	*Prob. to Switch*
Euronext Brussels	BEL20	98.08	1.92	95.36	4.64
Bulgarian Stock Exchange	SOFIX	97.17	2.83	89.15	10.85
Prague Stock Exchange	PX	98.92	1.08	95.28	4.72
Estonian Stock Exchange	OMXT	96.82	3.18	91.03	8.97
Helsinki Stock Exchange	HEX	98.96	1.04	97.34	2.66
Euronext Paris	CAC	98.24	1.76	96.04	3.96
Deustche Boerse	DAX	98.78	1.22	96.22	3.78
Athens Exchange, Athens Composite Index	ATH	98.15	1.85	95.82	4.18
Budapest Stock Exchange	BUX	98.61	1.39	95.92	4.08
Borsa Italiano	FMIB	99.06	0.94	96.79	3.21
Latvian Stock Exchange, Riga All-shares Index	OMXR	97.04	2.96	89.03	10.97
Vilnius Stock Exchange	OMXV	96.21	3.79	88.12	11.88
Amsterdam Stock Exchange	AEX	98.90	1.10	96.72	3.28
Oslo Stock Exchange	OSEAX	98.96	1.04	96.87	3.13
Warsaw Stock Exchange, 20	WIG20	99.18	0.82	97.29	2.71
Euronext Lisbon	PSI20	98.04	1.96	94.20	5.80
Bucharest Stock Exchange	BET	97.44	2.56	93.00	7.00
Moscow Exchange	MOEX	99.03	0.97	96.05	3.95
Bolsa de Madrid	IBEX	98.92	1.08	97.35	2.65
Stockholm Stock Exchange	OMXS	98.78	1.22	97.22	2.78
Swiss Exchange	SMI	98.92	1.08	95.79	4.21
Istanbul Stock Exchange	XU100	98.18	1.82	92.70	7.30
Ukraine Stock Exchange	UX	97.03	2.97	89.77	10.23
London Stock Exchange, FTSE 100	UKX	98.68	1.32	96.56	3.44

Source: Author's calculations.

A final insight from the Markov modelling is the remarkable persistence of regimes in European stock markets. If the index finds itself in bullish territory, the likelihood that it remains in it is in the range of 96% to 99%, leaving only a very small chance that a switch to a bearish market will occur. Similarly, once in a bearish market, an index has an average chance of 95% to remain bearish and only 5% to switch to the other regimes. Those averages naturally conceal individual country differences. Stock markets in countries like Lithuania, Latvia, Bulgaria and Ukraine have a much higher probability of switching in the range of 10-11%, while markets such as the Helsinki, Stockholm, Warsaw and Madrid stock exchanges register up to five times lower probability to do so. The high degree of inertia opens clear possibilities for risk hedging – once the HMM has successfully identified the current state of the market, it is highly likely that this will persist over the next period.

Figure 5. Regimes across Selected European Stock Markets: bullish (white) and bearish (shaded)

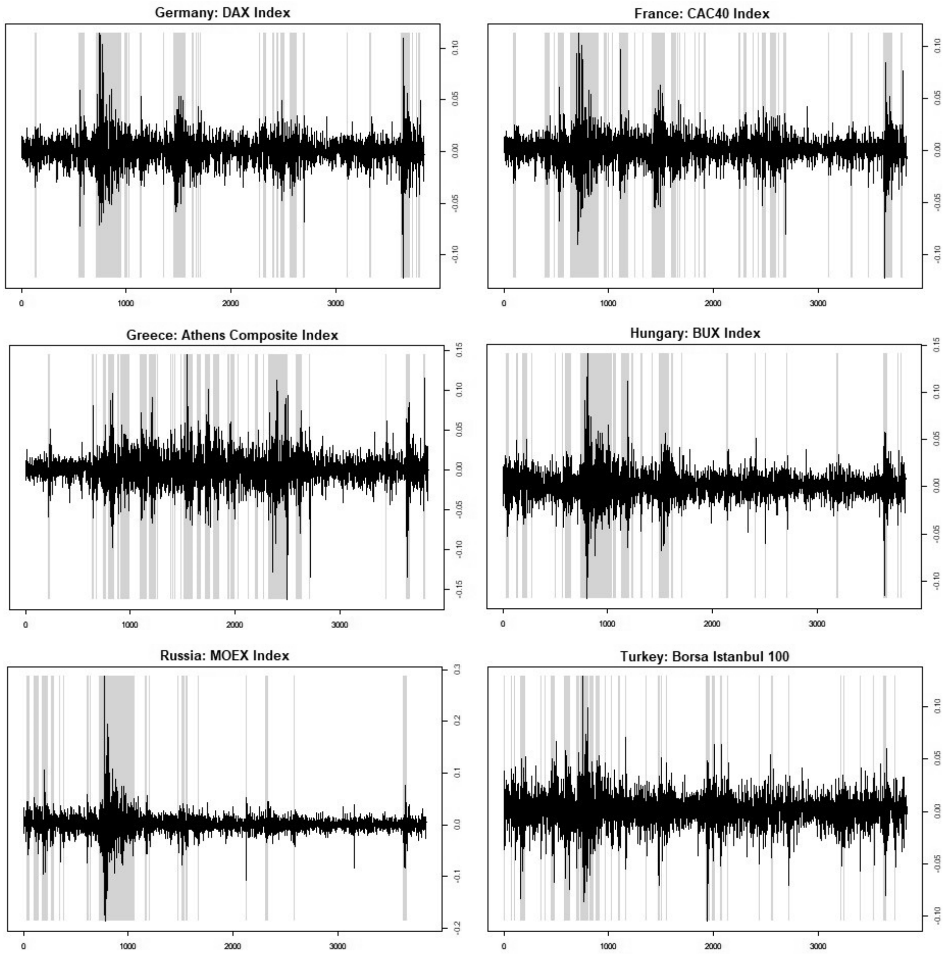

Source: Stooq Database and author's calculations.

A final pertinent point about risk diversification in European stock markets is the synchronization of different states. Should states be perfectly synchronized, risk reduction through capital transfers to an alternative stock exchange would not be possible. Conversely, the greatest possible degree of diversification can be obtained in two markets that have exactly opposite regime sequence. Figure 5 shows the synchronization pattern of six selected indices, whereby we observe imperfect correlation of state sequences. All markets are affected by the global financial crisis that began in 2008 in Europe and they simultaneously switch from a positive-return low-volatility bullish state to a high-volatility and low-return bullish ones. In this episode, synchronicity across markets is always perfect, thus making this episode an instance of non-diversifiable risk.

On the other hand, there are also many idiosyncratic regime switches across markets whereby some stock indices experience high volatility, while others remain in the calmer regime. Additionally, even when the overall regime sequences may coincide across markets, there is rarely perfect coincidence on a daily basis. For example, while the German and the French indices are roughly correlated, the regime switching does not necessarily occur on the same date, and this is particularly obvious in data from the turbulent 2020. Essentially this means that capital may find a safe haven moving away from markets with highly persistent bearish states into ones with highly persistent current bullish states.

This exercise in risk management has to be based on a daily analysis of current regimes and transition probabilities. Figure 5 also gives a new perspective on insights about returns correlations. Market indices that tend to be negatively correlated, such as the Istanbul stock exchange index on the one hand, and the CAC40 or Athens Composite, on the other, also display a significantly different regime structure. It is the change of the underlying market state that produces different returns and volatilities, which then result into overall negative correlations between markets. The study of regime switching thus has not only practical risk diversification and portfolio management implications but can also enable a deeper understanding of fundamental market functioning.

6. Implications and Conclusion

The selection and estimation of Hidden Markov Models to 24 European stock market indices revealed a number of insights. First, it supports the common wisdom and research results that markets are characterized by a limited number of very different states – e.g., bear and bull markets. The HMM selection made here is predicated on information criteria that clearly show the superiority of models with a small number of two or three states. Second, those states have very different characteristics, with one of them (bullish) being characterized by positive returns and low volatility, while the other one (bullish) – by overwhelmingly negative returns and a much higher level of fluctuations. Third, those states are remarkably persistent, with an average probability of switching from the calmer to the riskier state of less than 2%. However, when the switch does occur, the mode of market functioning is truly different and easy to detect. Fourth, the regime-switching is not perfectly synchronized across European markets and while there may be identical patterns, there are also many idiosyncratic episodes.

The results presented constitute a necessary but not sufficient condition for diversification and active management of an investment portfolio in European stocks. On the one hand, HMM regime estimation gives a clear indication of the risk-return tradeoff in all different markets under their two modes of operation, thus enabling the portfolio manager to select an optimal market for the current period, given their constraints. Detection of the current regime enables better estimation of expected returns given the current mode of operation and the expected volatility to be encountered. Transition matrices, on the other hand, provide the investor with a more sophisticated view of market persistence. Finally, leveraging regime detection and regime switching detection allows the coordinated movement of capital from markets with undesirable current regimes to those with more favourable ones. As European stock indices do not seem to be perfectly coordinated, this holds the promise of enabling further diversification and enabling better results than the strategy of simply taking and passively holding positions across markets.

On the other hand, investors will need to bear in mind not simply the asynchronous nature of stock markets but also other relevant characteristics that may need to be present in order for actual diversification opportunities to exist in practice. Among those, the liquidity, depth and breadth of markets, and the ensuing transaction costs and tax considerations readily come to mind. It may thus be of more interest of large institutional investors to focus on regime divergence between sizeable and liquid markets such as the stock markets in Germany, France, Italy, and the UK. Those markets will likely provide the necessary volume so that portfolio managers can easily reap the benefits of diversification. Conversely, small and illiquid markets may be of interest to smaller niche investors that require less market depth and may be under more liberal regulatory regimes. All in all, the results from the HMM should be considered as one of many tools of modern portfolio selection and management.

This paper has demonstrated how the application of Hidden Markov Models to European stock market indices enables the analyst not only to better estimate key market parameters and manage risk but also to gain a deeper understanding of the fundamental modes of operation of those stock markets. Even a simple and parsimonious Markov model with two regimes are able to succinctly capture different market dynamics and enable the study of the deeper drivers of those two distinct states. Markov models also reveal the structure of market correlations at a very granular level by showing state and regime switching patterns. This opens fruitful venues to study how markets are interconnected and what causes the lags between regime switching in different countries and stock markets. Further research into HMMs is needed to fully integrate this tool into the asset price forecasting and portfolio selection process and some progress is already being made. Work remains to be done in expanding the scope of application in terms of both investment tasks and alternative markets and investments, as well as conducting studies at the lower level of granularity of specific assets. Still, existing results indicate the significant utility of HMMs for both research and practical endeavours.

References

Acula, D. D., De Guzman, T. (2020). Application of Enhanced Hidden Markov Model in Stock Price Prediction. – Journal of Modeling and Simulation of Materials, 3(1), pp. 70-78.

– Economic Studies Journal (Ikonomicheski Izsledvania), 32(1), pp. 18-35.

Baltas, N., Karyampas, D. (2018). Forecasting the equity risk premium: The importance of regime-dependent evaluation. – Journal of Financial Markets, 38, pp. 83-102.

Cao, Y., Li, Y., Coleman, S., Belatreche, A., McGinnity, T. M. (2013). A hidden Markov model with abnormal states for detecting stock price manipulation. – In: 2013 IEEE International Conference on Systems, Man, and Cybernetics, pp. 3014-3019.

Cao, Y., Li, Y., Coleman, S., Belatreche, A., McGinnity, T. M. (2014). Adaptive hidden Markov model with anomaly states for price manipulation detection. – IEEE transactions on neural networks and learning systems, 26(2), pp. 318-330.

Damodaran, A. (2020). Equity Risk Premiums: Determinants, Estimation and Implications – The 2020 Edition. – SSRN Working Papers, 3550293.

De Angelis, L., Paas, L. J. (2013). A dynamic analysis of stock markets using a hidden Markov model. – Journal of Applied Statistics, 40(8), pp. 1682-1700.

Dias, J. G., Vermunt, J. K., Ramos, S. (2015). Clustering financial time series: New insights from an extended hidden Markov model. – European Journal of Operational Research, 243(3), pp. 852-864.

Goldfeld, S. M., Quandt, R. E. (1973). A Markov model for switching regressions. – Journal of Econometrics, 1(1), pp. 3-15.

Gupta, A., Dhingra, B. (2012, March). Stock market prediction using hidden Markov models. – In: 2012 Students Conference on Engineering and Systems, IEEE, pp. 1-4.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. – Econometrica: Journal of the econometric society, pp. 357-384.

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

Hassan, M. R. (2009). A combination of hidden Markov model and fuzzy model for stock market forecasting. – Neurocomputing, 72(16-18), pp. 3439-3446.

Hassan, M. R., Nath, B. (2005). Stock market forecasting using hidden Markov model: a new approach. – In: 5th International Conference on Intelligent Systems Design and Applications (ISDA'05), IEEE, pp. 192-196.

Kim, E. C., Jeong, H. W., Lee, N. Y. (2019). Global Asset Allocation Strategy Using a Hidden Markov Model. – Journal of Risk and Financial Management, 12(4), p. 168.

Kole, E. (2019). Markov Switching Models: An Example for a Stock Market Index. SSRN Working Paper 3398954.

Mizrach, B., Watkins, J. (1999). A Markov switching cookbook. – In: Nonlinear Time Series Analysis of Economic and Financial Data, Springer, Boston, MA, pp. 33-43.

Nguyen, N. (2017). An analysis and implementation of the hidden Markov model to technology stock prediction. – Risks, 5(4), p. 62.

Nguyen, N. (2018). Hidden Markov model for stock trading. – International Journal of Financial Studies, 6(2), p. 36.

Nguyen, N., Nguyen, D. (2015). Hidden Markov model for stock selection. – Risks, 3(4), pp. 455-473.

Pohle, J., Langrock, R., van Beest, F. M., Schmidt, N. M. (2017). Selecting the number of states in hidden Markov models: pragmatic solutions illustrated using animal movement. – Journal of Agricultural, Biological and Environmental Statistics, 22(3), pp. 270-293.

Rydén, T., Teräsvirta, T., Åsbrink, S. (1998). Stylized facts of daily return series and the hidden Markov model. – Journal of applied econometrics, 13(3), pp. 217-244.

Sanchez-Espigares, J. A., Lopez-Moreno, A., Sanchez-Espigares, M. J. A. (2015). Package 'MSwM'. traffic, 15, 1.

Sasikumar, R., Abdullah, A. S. (2016). Forecasting the stock market values using hidden Markov model. – International Journal of Business Analytics and Intelligence, 4, pp. 17-21.

Siegel, L. B. (2017). The equity risk premium: A contextual literature review. US: CFA Institute Research Foundation.

Somani, P., Talele, S., Sawant, S. (2014). Stock market prediction using hidden Markov model. – In: 2014 IEEE 7th Joint International Information Technology and Artificial Intelligence Conference, pp. 89-92.

Song, Y., Shi, S. (2012). Identifying speculative bubbles with an infinite hidden Markov model. The Rimini Center for Economic Analysis, Australia.

Sutkatti, R., Torse, D. (2019). Stock market forecasting techniques: A survey. – International Research Journal of Engineering and Technology, 6(5).

Zhang, X., Li, Y., Wang, S., Fang, B., Philip, S. Y. (2019). Enhancing stock market prediction with extended coupled hidden Markov model over multi-sourced data. – Knowledge and Information Systems, 61(2), pp. 1071-1090.