International Journal of Forecasting 39 (2023) 587–605

Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

Predictability of bull and bear markets: A new look at forecasting stock market regimes (and returns) in the US[☆]

Felix Haase [a], Matthias Neuenkirch [a,b,*]

[a] University of Trier, Germany
[b] CESifo, Germany

ARTICLE INFO

Keywords:
Forecast combination
Markov-switching models
Shrinkage methods
Stock market regimes
Time-varying transition probabilities

ABSTRACT

The empirical literature of stock market predictability mainly suffers from model uncertainty and parameter instability. To meet this challenge, we propose a novel approach that combines dimensionality reduction, regime-switching models, and forecast combination to predict excess returns on the S&P 500. First, we aggregate the weekly information of 146 popular macroeconomic and financial variables using different principal component analysis techniques. Second, we estimate Markov-switching models with time-varying transition probabilities using the principal components as predictors. Third, we pool the models in forecast clusters to hedge against model risk and to evaluate the usefulness of different specifications. Our weekly forecasts respond to regime changes in a timely manner to participate in recoveries or to prevent losses. This is also reflected in an improvement of risk-adjusted performance measures as compared to several benchmarks. However, when considering stock market returns, our forecasts do not outperform common benchmarks. Nevertheless, they do add statistical and, in particular, economic value during recessions or in declining markets.

1. Introduction

The existence of different stock market regimes is widely accepted among academics and practitioners. Stock market cycles typically precede business cycles and are caused by time-varying expectations of future cash flows and discount rates. In bullish periods, prices rise and fluctuate only mildly, whereas in bearish periods, prices decrease and volatility increases. Hence, anticipating regime changes and, in particular, contractions is of relevance for investors and corporate decision-makers. Furthermore, the state of the stock market as leading indicator is important for governments, (central) banks, and households. The global financial crisis (GFC) of 2007–2008 is the most recent example illustrating the danger of spillover effects to the real economy.

Since stock market regimes are unobservable, their identification and prediction is challenging. Three methods have been established in the literature. First, observable measures that reflect the risk aversion of market participants are natural candidates to signal regime dynamics. Empirically, Coudert and Gex (2008) highlight the relevance of risk-aversion proxies for stock crash predictions, whereas Chow et al. (1999) and Kritzman and Li (2010) underline the importance of market turbulence indices. Second, Markov-switching (MS) models are used to infer the probabilities of a latent state variable and to forecast returns or volatility (Ang and Bekaert 2002; Haas et al. 2004); the number of regimes in these models is still subject to debate (e.g., Guidolin and Timmermann 2007; Maheu et al. 2012; Hauptmann et al. 2014). Third,

☆ We would like to thank seminar participants at Freie Universität Berlin, the Associate Editor, and two anonymous reviewers for their helpful comments on an earlier version of the paper. The initial draft of the paper circulated under the title "Forecasting Stock Market Recessions in the US: Predictive Modeling using Different Identification Approaches".
* Correspondence to: University of Trier, Department of Economics, D-54286 Trier, Germany.
E-mail address: neuenkirch@uni-trier.de (M. Neuenkirch).

https://doi.org/10.1016/j.ijforecast.2022.01.004
0169-2070/

change point detection methods or dating rules are utilized in this context. The application of change point analysis to stock market data is similar to MS models (Pástor and Stambaugh 2001; Pettenuzzo and Timmermann 2011). However, the assumption that "history repeats" is neglected, so that each change point marks the beginning of a new regime. Dating rules search for local extremes which are defined by period lengths (Pagan and Sossounov 2003) or by absolute price changes (Lunde and Timmermann 2004). The underlying algorithms need past and future prices for the dating of recessions, and, consequently, delayed signals may occur. In addition, genuine backtesting cannot be performed in real time when using dating rules.

Considering the empirical success of dimensionality reduction techniques (Çakmaklı and van Dijk 2016, Neely et al. 2014), regime-switching models (Guidolin and Hyde 2012, Maheu et al. 2012), and forecast combination (Rapach et al. 2010) in predicting stock market dynamics, we propose a novel procedure that combines these approaches. Confronted with a large real-time dataset of macroeconomic and financial market variables, we first reduce the dimensionality into a few latent factors by different principal component analysis (PCA) techniques. We employ a conventional PCA and a sparse PCA, where the loadings of some variables are set to zero. In addition, we apply a soft thresholding approach to both conventional PCA and sparse PCA, yielding two additional sets of targeted principal components. Second, using the principal components as predictors, we estimate MS models with time-varying transition probabilities (TVTP) to identify and predict regimes in a single step. For this purpose, we consider two specifications. On the one hand, we use a general specification, which models the conditional mean and the transitions (Specification A). And on the other hand, we rely on a restricted specification where only regimes are predictable while returns follow a (regime-dependent) random walk (Specification B). Since highly parameterized models tend to be inferior to parsimonious ones in terms of forecast accuracy, we limit the model size of each model to include only one principal component (or observable predictor). These different combinations of MS specifications and PCA techniques (or the usage of observable predictors) result in a large number of models that we combine into several forecast clusters (according to the shrinkage method and the model specification). In this third step, we also ensure robustness to different weighting decisions as we consider simple averaging and a continuous weighting approach. Throughout the procedure, we account for publication lags, data revisions, and consider transaction costs to ensure realistic forecasts in the backtest. Fig. 1 (at the beginning of Section 2) provides an illustrative overview of our methodology, and Table 1 (at the end of Section 2) summarizes the different specifications, clusters, models, and forecast combination techniques.

Our sample consists of weekly data for the S&P 500 and spans the period from November 17, 1989 to May 7, 2021. We use weekly data since, at a higher frequency, regime forecasts would be too noisy and return forecasts virtually impossible. Moreover, the choice of weekly returns represents a good compromise between precision and data availability, as fundamentals are usually updated monthly, while market data obviously change at an intraday frequency. Our recursive out-of-sample real-time exercise focuses on the most recent 864 weeks. Accordingly, the first training set to estimate the MS models ends on October 15, 2004. For the evaluation of the different forecasts, we classify bull and bear markets using the dating rule of Lunde and Timmermann (2004), assuming knowledge of the full sample.

Our regime forecasts are suitable to respond to regime changes in a timely manner to participate in recoveries or to prevent losses. This is also reflected in an actual economic value added, as many of our forecasts beat all benchmarks in risk-adjusted performance measures. However, when considering stock market returns, our forecasts do not statistically outperform common benchmarks. The fact that return forecasts perform worse than regime forecasts is not surprising, since forecasting the broader trend of the stock market is obviously easier than providing point forecasts, in particular at a weekly frequency. Nevertheless, our return forecasts still provide some economic value added for risk-averse investors, as they generate a lower annualized standard deviation of the returns and better tail-risk measures than the corresponding regime forecasts. Consistent with the literature (e.g., Henkel et al. 2011, Rapach and Zhou 2013), we find that much of the predictability comes from periods of market turmoil or recessions. Finally, we highlight that it is sufficient to model the time-varying conditional transitions in a Markov-switching model. We also suggest relying on dimensionality reduction techniques and enhancing the conventional principal component analysis with shrinkage methods such as sparsity and/or soft thresholding.

Our paper is, to the best of our knowledge, the first one to apply MS models with TVTP and several PCA techniques to predict bull and bear markets. We contribute to several strands of the stock market forecasting literature. First, we confirm the previous finding of predictable trends, in particular during recessions, in stock markets (Chen 2009, Guidolin and Timmermann 2007, Kritzman et al. 2012).

Second, we emphasize the benefits of MS models with principal components and TVTP. Although MS models with time-varying transitions were developed more than 25 years ago (Diebold et al. 1994), there are only a handful of examples that apply these models in the context of bull and bear markets (e.g., Focardi et al. 2019, Guidolin and Hyde 2012, Kole and van Dijk 2017, Maheu and McCurdy 2000, Schaller and Norden 1997). The few existing papers that include macro-financial variables in the transition equation provide rather disappointing results. Guidolin and Hyde (2012) and Kole and van Dijk (2017) do not detect any advantage of modeling the transition with lagged returns, individual macro-financial variables, or a principal component based on seven popular predictors. Overly complex modeling of the switching process might cause their results, since Guidolin and Hyde (2012) apply a multivariate three-regime model, and Kole and van Dijk (2017) consider multiple variables in the switching equation. We address these concerns in two ways. On the

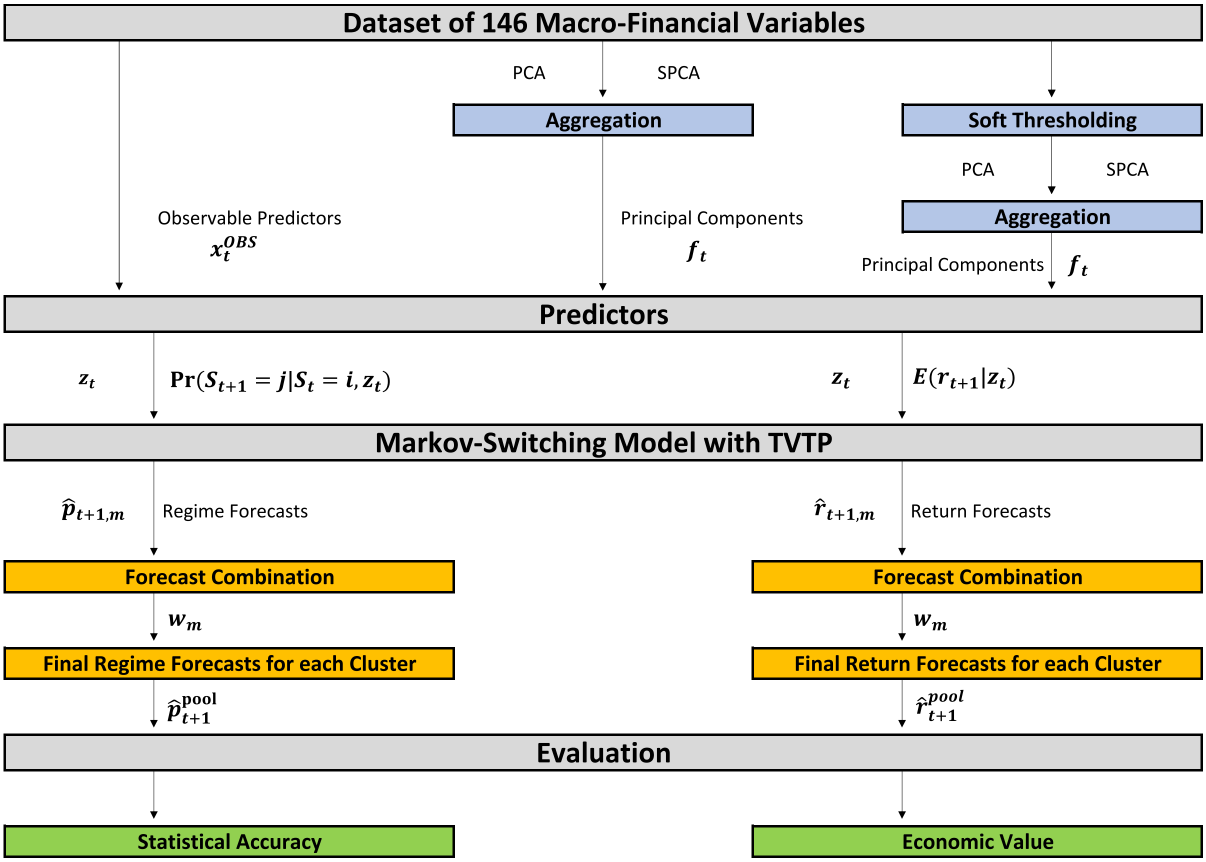

Fig. 1. Overview of the methodology.

one hand, we follow the recommendation of Zens and Böck (2019) to include only a few latent factors (one, to be precise) into the transition equation and, on the other hand, we focus on a univariate setting with only two regimes to reduce possible identification problems and estimation uncertainty.

Third, for constructing stock predictors from "big data", we recommend using shrinkage methods. These provide a more straightforward interpretation of the extracted factors and, particularly appealing in forecasting, can reduce noise without losing much of the captured variance. This finding is also documented by Rapach and Zhou (2019), who emphasize the superiority of a sparse PCA. Finally, there are several other papers that use a high-dimensional dataset to predict financial variables. Mönch (2008) and Ludvigson and Ng (2009) predict bond yields. Ludvigson and Ng (2007), Neely et al. (2014), and Çakmaklı and van Dijk (2016) provide promising results and highlight the attractiveness of principal components as predictors for stock returns. Our empirical framework also resembles the ones used to forecast commodity returns and futures with large datasets and regime-switching models (Guidolin and Pedio 2020, 2021) or to exploit the business cycle to predict asset prices (e.g., Hammerschmid and Lohre 2018, Kaya et al. 2010, Sander 2018).

The remainder of this paper is organized as follows. Section 2 outlines our methodology and explains the necessary modeling choices. Section 3 introduces the dataset of macro-financial variables. Section 4 shows the classification of market regimes and discusses the aggregation of the predictors, both under the assumption of full-sample knowledge. Section 5 demonstrates how our approach works in a real-time situation with recursive out-of-sample forecasts. Section 6 concludes.

2. Methodology

We face the issues of model uncertainty and parameter instability when forecasting stock market regimes and returns (Pesaran and Timmermann 1995). Our approach combines dimensionality reduction, regime-switching models, and forecast combination to predict excess returns on the S&P 500. In addition, we apply the MS specifications and the forecast combination schemes to a subset of directly observable popular predictors. Our aim is to evaluate whether a large dataset and the utilization of aggregation techniques (see, among others, Neely et al. 2014 and Çakmaklı and van Dijk 2016) provide an actual advantage over employing commonly used (simple) predictors.[1] Fig. 1 provides an overview of the individual steps in our procedure that are explained in detail in the following subsections.

[1] We consider the lagged return R, the dividend-price ratio DP (Campbell and Shiller 1988, Fama and French 1988, Schaller and Norden 1997), the volatility index VIX (Rubbaniy et al. 2014), the term spread TS and the credit spread CS (Campbell and Yogo 2006, Fama

2.1. Step 1: Data aggregation

Due to the increasing availability of data, an investor is confronted with the choice of the relevant predictors. Theoretical considerations might be helpful in this context, but even with certain restrictions there is a large pool of potential variables. Due to the substantial correlation of many covariates with unobserved state variables—such as the business cycle or investor sentiment—an efficient filtration of the variables is recommended to cover the comovement and to eliminate potential noise. PCA is an appealing method to capture relevant information in a parsimonious way. A small number of components is usually sufficient to capture most of the variation in the data, allowing for a significant reduction in the dimensionality of the original dataset.

2.1.1. Conventional PCA and sparse PCA

Conventional PCA: Principal components capture the comovement of many (potentially) correlated predictors that are normalized to a mean of zero and a variance of one. Let X be a $T \times K$ matrix of potential predictors, where the number of rows T $(t = 1, 2, \ldots, T)$ represents the time dimension and K $(k = 1, 2, \ldots, K)$ the cross-sectional dimension. Using singular value decomposition of X, we can obtain the principal components as (Zou et al. 2006):

$$X = UDV^T \tag{1}$$

The principal components are $Z = UD$, with U representing a unitary matrix and D a diagonal matrix of singular values. V is a $K \times K$ matrix of eigenvectors, where the kth column represents the loadings of the kth component. Typically, a small positive number of q components is sufficient to aggregate the information in X, so that we achieve a substantial dimensionality reduction in exchange for a minimal loss of information ($q \ll min(K, T)$). In addition, the components are constructed to be uncorrelated to each other. To determine q, we use the IC_{p2} information criterion by Bai and Ng (2002), where the upper bound is set according to an automatic elbow procedure. Hence, we select the first q normalized principal components as relevant factors f to predict stock market regimes and returns.

Sparse PCA: One disadvantage of conventional PCA is that the components are based on all variables, which often leads to a lack of interpretability. A sparse PCA uses shrinkage methods to reduce the loadings of some variables to zero for a more straightforward interpretation and a lower risk of overfitting without losing too much of the captured variance (Rapach and Zhou 2019). Following the illustration of Zou et al. (2006), we treat the optimization as a regularized regression problem. Suppose we consider the first q principal components, and let x_t be the tth row of X. We further denote A as a $q \times K$ orthonormal matrix with elements $A = [\alpha_1, \alpha_2, \ldots, \alpha_K]$ and B as a $q \times K$ sparse weight matrix with $B = [\beta_1, \beta_2, \ldots, \beta_K]$. Then we consider the following optimization problem for $\lambda > 0$:

$$\underset{A,B}{\arg\min} \left[\sum_{t=1}^{T} \|x_t - AB^T x_t\|^2 + \lambda \sum_{p=1}^{q} \|\beta_p\|_2^2 + \sum_{p=1}^{q} \lambda_{1,p} \|\beta_p\|_1 \right] \quad s.t.\ A^T A = I \tag{2}$$

$\|.\|_1$ corresponds to the $L1$ and $\|.\|_2^2$ to the squared $L2$ norm. I represents the $q \times q$ identity matrix. The amount of ridge shrinkage λ is the same for all q components and the sparsity constraint $\lambda_{1,p}$ can vary over the components, where a higher value of $\lambda_{1,p}$ leads to more sparse loadings. If we restrict Eq. (2) by $B = A$ and set the LASSO (least absolute shrinkage and selection operator) penalty $\lambda_{1,p} = 0$, we obtain the conventional PCA (Zou et al. 2006). We solve Eq. (2) using the variable projection approach by Erichson et al. (2020).

For a better comparability, we do not apply the procedure of Bai and Ng (2002) on the adjusted sparse factors (see Zou et al. 2006). Instead, we assume the same number of components as the conventional PCA suggests. Additionally, we do not vary the degree of sparseness over the components and set $\lambda_{1,p} = \lambda_1$. For the $L1$ and $L2$ penalty, we follow Kristensen (2017) and tune the hyperparameters in every week such that the Bayesian information criterion (BIC) is minimized.[2]

2.1.2. Soft thresholding

Another drawback of the conventional PCA and the sparse PCA is that they do not consider the target variable during the construction of the factors. A soft thresholding approach (Bai and Ng 2008) conducts a preselection on the data to obtain targeted predictors and has already been applied in the return forecasting literature (e.g., Çakmaklı and van Dijk 2016). Our implementation of soft thresholding follows Bai and Ng (2008) and uses the elastic net (EN) methodology. The EN is a convex combination of LASSO and ridge regression that performs model selection and shrinkage simultaneously.[3] More formally, the EN optimization is a regularized regression to minimize the residual sum of squares (RSS) and can be written as

and French 1989), the Purchasing Managers Index *PMI* (Johnson and Watson 2011), and the variance risk premium *VP* (Bekaert and Hoerova 2014, Bollerslev et al. 2009). The term spread is defined as difference between the 10 Y US Treasury bond and the 3 M Treasury bill, the credit spread as excess yield of the Moody's seasoned Baa over the Aaa corporate bond yield, and the variance risk premium as the difference between the squared *VIX* and the sum of the squared 5-minute returns of the last 22 trading days.

[2] For this purpose we use the grids $\lambda \in (1e^{-4}, 1e^{-3}, 1e^{-2})$ and $\lambda_1 \in (2e^{-3}, 4e^{-3}, 6e^{-3}, 8e^{-3}, 1e^{-2})$:

$$(\lambda_1^*, \lambda^*) = \underset{\lambda_1,\lambda}{\arg\min}\, log \left(\frac{1}{KT} \sum_{k=1}^{K} \sum_{t=1}^{T} [x_{k,t} - \hat{l}_{SPCA,k}(\lambda_1, \lambda) \hat{F}_{SPCA,t}(\lambda_1, \lambda)]^2 \right) + \phi(\lambda_1, \lambda) \frac{log(KT)}{KT} \tag{3}$$

$\phi(\lambda_1, \lambda)$ represents the number of non-zero PCA weights, and $\hat{l}_{SPCA}$ and $\hat{F}_{SPCA}$ correspond to the loadings and adjusted scores (via QR decomposition), as suggested by Zou et al. (2006).

[3] The main benefit of EN over LASSO in soft thresholding is that in situations with a group of highly correlated predictors, LASSO selects only one variable of this group, whereas the EN approach stretches "the fishing net to retain all the *big fish*" (Bai and Ng 2008, p. 307).

follows:

$$\arg\min_{\beta}\left[RSS+\lambda_1\sum_{k=1}^{K}|\beta_k|+\lambda_2\sum_{k=1}^{K}\beta_k^2\right] \quad (4)$$

β corresponds to the EN estimate and λ_1 and λ_2 are non-negative hyperparameters, which balance the influence of LASSO and the ridge penalty. In our context, we use the non-zero β to select relevant predictors.

The choice of the target variable depends on our objective, resulting in different sets of predictors. If we want to find targeted predictors for the return process, we rely on future excess returns r_{t+1}. However, if we want to select predictors to forecast regimes, our target cannot be observed. Here, we proceed with the VIX_{t+1}, which is a popular fear gauge in practice and, therefore, a good signal for shifts into a bearish regime. We follow Bai and Ng (2008) and use the least angle regression algorithm to solve the elastic net problem (LARS-EN). We obtain a ranking of selected predictors, such that we can substitute the LASSO penalty λ_1 with the size of the active set of predictors. We refrain from optimizing the size of the active set for simplicity and select the top 75 predictors, which is proportionally similar to the subset size in Çakmaklı and van Dijk (2016). With respect to the ridge penalty, we perform a grid search on the interval [0, 0.25, 0.5, 0.75, 1] and choose the λ_2^* that optimizes Mallows's Cp. For the out-of-sample exercise, we repeat the hyperparameter search every week.

We apply the soft thresholding approach in combination with both conventional PCA and sparse PCA. Hence, as predictors, we utilize four different sets of principal components and, additionally, the subset of directly observable popular variables (see Fig. 1).

2.2. Step 2: Markov-switching models

Since the pioneering work of Hamilton (1989), MS models have become increasingly popular in economics. MS models are able to reveal changes in the fundamental environment of financial markets in a timely manner, even if their interpretation is only possible ex post (Ang and Timmermann 2012). Thus, MS models help to account for time-varying risk premia and to uncover temporary trends in returns.

Starting with the basic switching model, r_t denotes the excess log-return of the S&P 500 over the 3 M Treasury bill and S_t the unobservable state of the stock market. Then, the non-linear return dynamics can be described as:

$$\begin{aligned} r_t &= \mu_{S_t} + u_t \\ u_t &\sim i.i.d.\ N(0,\sigma_{S_t}^2) \\ Pr(S_t = j|S_{t-1}=i) &= p_{ij} \end{aligned} \quad (5)$$

Assuming that the mean μ_{S_t} and the variance $\sigma_{S_t}^2$ are dependent on the current market regime, the MS model is able to replicate stylized facts of financial time series such as fat tails, volatility clustering, and asymmetries (Ang and Timmermann 2012). In the basic time-homogeneous case, the regime variable S_t is assumed to follow a discrete first-order Markov chain. That is, the current market regime j depends only on the previous regime i. We refer to this model, which serves as one of the benchmarks below, as an MS model with time-constant transition probability (TCTP) and without external predictors.

The majority of papers treat the transition probabilities as constant over time, ignoring that these can be affected by changes in fundamental conditions. In this paper, we follow Diebold et al. (1994) and model the switching process as being dependent on macro-financial conditions z_{t-1}.

Specification A: In the general specification, we assume that the excess S&P 500 returns follow an MS model with predictable mean and regime processes:

$$\begin{aligned} r_t &= \mu_{S_t} + \beta_{S_t} z_{t-1} + u_t \\ u_t &\sim i.i.d.\ N(0,\sigma_{S_t}^2) \\ p_{i0,t} &= \frac{\exp(\upsilon_{i0}+\gamma_{i0}z_{t-1})}{1+\exp(\upsilon_{i0}+\gamma_{i0}z_{t-1})} \end{aligned} \quad (6)$$

z_{t-1} are either observable predictors proposed by the literature or principal components obtained using the different techniques described in the previous subsection. The intercept is denoted as μ_{S_t}, and u_t is the idiosyncratic error with a regime-dependent variance. To model the switching dynamics, we follow the standard in the literature by using a logit link function (Diebold et al. 1994), where the constant υ_{i0} and the slope γ_{i0} depend on the current regime. Finally, it has to be noted that all parameters are dependent on the regime variable S_t, allowing for parameter flexibility across regimes.

In our application, we consider two regimes, where regime 0 corresponds to bull markets and regime 1 to bear markets. The number of stock market regimes is certainly open to debate, and, since S_t is a latent variable, the "true" number is unknown. An approximation with econometric tests is also difficult (Ang and Timmermann 2012, Hansen 1991). Therefore, one usually relies on information criteria or theoretical arguments. Our decision to focus on two regimes is motivated by several reasons. First, a clear distinction can be made between (i) a volatile regime with a negative drift and (ii) a calm regime with positive average returns. Second, prominent dating rules (Lunde and Timmermann 2004, Pagan and Sossounov 2003) are available for two regimes. These ensure a transparent and straightforward regime classification and are helpful to evaluate our real-time regime predictions ex post. Finally, more than two regimes often lead to unstable estimations, particularly in our out-of-sample task with a variety of predictors and specifications.[4]

As highlighted by Zens and Böck (2019), only a small number of variables can be included in the transition probabilities to ensure a stable estimation process. Consequently, we rely on latent factors constructed from many variables to incorporate macro-financial information in a compact form and restrict the number of variables to avoid highly parameterized models. More precisely, we

[4] Note that some authors assume more than two regimes (Guidolin and Hyde 2012, Guidolin and Timmermann 2007, Maheu et al. 2012, Zhu and Zhu 2013). For example, Maheu et al. (2012) distinguish between two bullish regimes (normal and correction) and two bearish regimes (normal and rally).

Table 1
Specifications, clusters, models, and forecast combinations.

Specification	Cluster	Models	Forecast combination	
Specification A: Conditional mean and transitions	OBS ($M = 7$)	A-R, A-DP, A-VIX, A-TS, A-CS, A-PMI, A-VP	A-OBS-AVE	A-OBS-BMA
	PC ($M^{max} = 8$)	A-PC1, A-PC2, . . ., A-PCq	A-PC-AVE	A-PC-BMA
	SPC ($M^{max} = 8$)	A-SPC1, A-SPC2, . . ., A-SPCq	A-SPC-AVE	A-SPC-BMA
	TPC ($M^{max} = 7$)	A-TPC1, A-TPC2, . . ., A-TPCq	A-TPC-AVE	A-TPC-BMA
	TSPC ($M^{max} = 7$)	A-TSPC1, A-TSPC2, . . ., A-TSPCq	A-TSPC-AVE	A-TSPC-BMA
Specification B: Transitions only	OBS ($M = 7$)	B-R, B-DP, B-VIX, B-TS, B-CS, B-PMI, B-VP	B-OBS-AVE	B-OBS-BMA
	PC ($M^{max} = 8$)	B-PC1, B-PC2, . . ., B-PCq	B-PC-AVE	B-PC-BMA
	SPC ($M^{max} = 8$)	B-SPC1, B-SPC2, . . ., B-SPCq	B-SPC-AVE	B-SPC-BMA
	TPC ($M^{max} = 7$)	B-TPC1, B-TPC2, . . ., B-TPCq	B-TPC-AVE	B-TPC-BMA
	TSPC ($M^{max} = 7$)	B-TSPC1, B-TSPC2, . . ., B-TSPCq	B-TSPC-AVE	B-TSPC-BMA

Notes: All MS models are estimated with TVTP. Specification A contains predictors in the switching equation and the conditional mean equation according to Eq. (6). Specification B contains predictors in the switching equation only. OBS: observable predictors; PC: conventional PCA; SPC: sparse PCA; TPC: conventional PCA with soft thresholding; TSPC: sparse PCA with soft thresholding; R: lagged return; DP: dividend-price ratio; TS: term spread; CS: credit spread; PMI: Purchasing Managers Index; VP: variance risk premium; AVE: simple average; BMA: Bayesian model averaging. For example, model A-SPC3 is estimated with the third sparse principal component as the conditional mean predictor and transition predictor z_{t-1}. Model B-R uses lagged returns as the transition predictor z_{t-1}. A-TSPC-AVE uses simple averaging to combine the forecasts from the cluster TSPC. The number of estimated principal components varies (i) for the full dataset of 146 macro-financial variables versus the targeted dataset of 75 variables and (ii) for the in-sample forecasts versus the recursive out-of-sample forecasts. See Section 4.2 and Appendix B.

incorporate only one principal component (or observable variable) in the switching equation and in the conditional mean equation. This ensures a robust estimation process and reduces the variability of the forecasts. For simplicity, we assume that the external predictors are the same in both equations.

Specification B: Given the extensive regime dependency of Specification A, overfitting might be a problem. For this reason, we also consider a restricted model that focuses only on the switching process while returns follow a (regime-dependent) random walk. By setting the constraint $\beta_{S_t} = 0$ in Specification A, we obtain the restricted Specification B.

We estimate all models with maximum likelihood methods using the expectation maximization algorithm.[5]

Prediction: One appealing feature of MS models is that identification and prediction can be done in a single step. Using the filter proposed by Hamilton (1989), the one-step-ahead regime prediction for j is:

$$\hat{p}^j_{t+1} = Pr(S_{t+1} = j|\Omega_t) = \sum_{i=0}^{1} p_{ij,t} Pr(S_t = i|\Omega_t) \tag{7}$$

Ω_t represents the information set in period t and $Pr(S_t = i|\Omega_t)$ the filtered probability, which is recursively updated using Bayes' rule. To simplify the notation, we define $\hat{p}^1_{t+1} = \hat{p}_{t+1}$ as the predicted bear probability and $(1 - \hat{p}_{t+1})$ as the corresponding bull probability.

[5] Hereby, we essentially follow Hamilton (1990). An alternative would be a Bayesian approach using the Gibbs sampler, in which parameter uncertainty is explicitly incorporated (for an application, see Maheu et al. 2012). For further details about inference on regimes and the estimation procedure, we refer the reader to Hamilton (1994).

Finally, the regime forecasts can be used to predict returns. Relying on the regime-dependent expectations $E[r_{t+1}|S_{t+1} = j]$, the return forecast $\hat{r}_{t+1}$ is given by the following probability-weighted average:

$$\hat{r}_{t+1} = (1 - \hat{p}_{t+1})E[r_{t+1}|S_{t+1} = 0] + \hat{p}_{t+1}E[r_{t+1}|S_{t+1} = 1] \tag{8}$$

2.3. Step 3: Forecast combination

Instead of using multiple predictors in one model, forecast combination uses multiple models with a restricted number of predictors in each model. Timmermann (2006) highlights that combined forecasts work particularly well in uncertain situations where the influence of relevant variables varies considerably over time. Hence, forecast combination is a promising strategy to hedge against model uncertainty and to increase the predictability of regimes (and returns). Compared to large multivariate regressions, forecast combination has the advantage that the estimation variability can be significantly reduced and that in-sample overfitting can be avoided (Rapach and Zhou 2013). In general, the forecast combination setting can be formulated as a weighted average of individual forecasts for regimes and returns.

Regime forecasts: Suppose we have M regime probability forecasts $\hat{p}_{t+1,m}$. This yields the following forecast combination problem:

$$\hat{p}^{pool}_{t+1} = \sum_{m=1}^{M} w_m \hat{p}_{t+1,m} \tag{9}$$

Return forecast: The pooled return forecast, given M return forecasts $\hat{r}_{t+1,m}$, can be expressed as follows:

$$\hat{r}^{Pool}_{t+1} = \sum_{m=1}^{M} w_m \hat{r}_{t+1,m} \quad (10)$$

In this context, we have to make a decision about the number of included forecasts M and their weights w_m. In our application, the individual forecasts are combined within some prespecified clusters. We form the clusters in such a way as to be able to evaluate the usefulness of the various aggregation techniques and the specification choices of the MS model. Consequently, we differentiate alongside two dimensions: (i) predictor choice (directly observable or estimated using the four different PCA techniques), and (ii) MS specification (Specification A or Specification B).

Next, we have to determine the individual weights of the forecasts w_m. For this purpose, we employ two different methods:

Simple average (AVE) $\quad w_m = \frac{1}{M}$

Bayesian model averaging (BMA) $\quad w_m = \frac{exp(-\Delta_m/2)}{\sum_{l=1}^{L} exp(-\Delta_l/2)}$

The simple average forecast is straightforward and precludes any estimation risk. In addition, it often provides good results, which are difficult to outperform (Timmermann 2006). In addition, inspired by the results of Cremers (2002), we apply Bayesian model averaging. Since our estimation is not Bayesian, we approximate the posterior model probability with the observed data. We use Bayes' factors to avoid computational difficulties (overflow/underflow) and define $\Delta_m = BIC_m - BIC^*$, where BIC^* represents the model with the lowest BIC.

To summarize, we calculate a total of 20 forecast combinations. This number emerges from the five clusters of predictors (observable predictors and the four different PCA techniques), the two specifications of the MS model (mean and transitions versus transitions only), and the two different aggregation techniques (simple average versus BMA). Table 1 summarizes the different specifications, clusters, models, and forecast combination techniques.

3. Data

Our dataset consists of weekly data for the United States. The stock market is represented by the S&P 500 index, adjusted for dividends and stock splits. We consider a large set of 146 variables to predict regimes and returns. This includes several categories of variables: bond yields, term spreads and credit spreads, lagged returns, technical indicators, industry returns, market-based risk indicators, valuation ratios, survey-based expectations about macroeconomic variables/earnings and their dispersion, sentiment indicators, and macroeconomic fundamentals. All variables either are proven to be empirically relevant or can be recommended from a practical point of view.

The bond market reflects expectations of market participants in terms of growth prospects, future interest rates, projected inflation, and current risk aversion. Among others, Estrella and Mishkin (1996, 1998) point out that information extracted from the yield curve and, in particular, term spreads are robust predictors for recessions in the real economy. Therefore, we consider government bond yields of all available maturities as well as various spreads over different maturities and the London Interbank Offered Rate (LIBOR). Since stock market contractions are often induced by an increase in risk aversion, credit spreads might also be useful in this context (Coudert and Gex 2008). Correspondingly, we take corporate bond spreads from Moody's and the TED (Treasury bill Eurodollar) spread into account. As additional predictors, we consider the realized variance of the S&P 500 expressed as sum over the 5-minute squared returns of the previous 1, 5, and 22 trading days plus the close-to-open return (see Bollerslev et al. 2009). Furthermore, we use information from option markets by using the implied volatility index of the S&P500, the VIX. Following Bollerslev et al. (2009), the VIX can be decomposed in a component that reflects the expected future volatility and a risk premium. We extract the so-called variance risk premium by subtracting the squared VIX from the realized stock market variance of the last 22 trading days. Finally, we use additional indicators that capture changes in risk perception, like the gold price and the West Texas Intermediate (WTI) oil price.

We also utilize survey-based expectations as predictors. Consensus Economics asks analysts from banks and research institutes about their macroeconomic expectations at monthly intervals. As predictors, we employ the first and second moments of the individual one-year-ahead expectations of macroeconomic variables and the three- and 12-month-ahead interest rate expectations. The macroeconomic expectations are complemented by sell-side analysts' earnings forecasts, their revisions, and their dispersion from the Institutional Brokers Estimate System. In addition, we employ sentiment measures, such as the surveys by the Conference Board. Following Chen (2012), we also consider several consumer confidence measures as predictors. To capture broader macroeconomic expectations, we utilize the leading composite index from the Conference Board and the PMI. Lastly, we roughly consider the same standard macroeconomic variables as Chen (2009) to incorporate previous findings into our analysis.[6]

The current valuation level is typically related to stock market turbulence (Campbell and Shiller 1988, Fama and French 1988, Lewellen 2004). Hence, we include the dividend-price ratio, the earnings-price ratio, the 10 Y earnings-price ratio, and the payout ratio in our dataset. Moreover, we use the same technical indicators as those proposed by Neely et al. (2014). In addition, we incorporate the short-run and long-run moving average of returns (1 M and 12 M) into our predictor set, which are either equally or exponentially weighted. It might be argued that price "excesses" are a major cause of future contractions, which suggests that valuation ratios or historical returns

[6] Industrial production, M1 and M2, the inflation rate, and the unemployment rate.

correlate positively with the risk of bear markets. Furthermore, signals from technical indicators are highly relevant in practice and reflect psychological aspects.

We also use the returns of 34 industry portfolios from the Center for Research in Security Prices Database. Hong et al. (2007) point out that the broad market often processes the information diffused in the industrial returns with a delay, which highlights the leading character of some industry returns. Additionally, we calculate the financial turbulence index (Chow et al. 1999, Kritzman et al. 2012) as well as the absorption ratio (Kritzman et al. 2011). Both measures are popular choices to detect anomalies. The financial turbulence index signals convergence and divergence regarding historical correlation structures and extreme price movements. The absorption ratio can be seen as proxy of systematic risk and encompasses the captured variance of a rolling PCA with a fixed number of components. Since this measure is relatively persistent, we rely on the standardized change in the absorption ratio. To calculate these two risk indicators, we follow the methodology of Kritzman et al. (2011, 2012).

Our sample spans the period between November 17, 1989 and May 7, 2021.[7] Our out-of-sample real-time exercise is conducted using the most recent 864 weeks. Correspondingly, the first training set to estimate the MS models ends on October 15, 2004. Starting from this date, we employ a recursive scheme with an expanding window to predict regimes and returns in the US. In all cases, we rely on end-of-week data (if the data are available at a higher frequency). Every variable is shifted to its publication date and we account for data revisions to ensure a real-time perspective. In addition, we apply common transformations to ensure stationary predictors as, for instance, we follow Rapach et al. (2005) and detrend bond yields by their one-year moving average. Finally, all variables are centered and scaled before their inclusion in the prediction model. Appendix A lists all variables alongside their definitions (Table A1), provides the data sources (Table A2), and displays summary statistics (Table A3).

4. In-sample results

The focus of this paper is on the out-of-sample performance of our forecasts. Hence, we keep the discussion of the in-sample results as concise as possible. Accordingly, we focus on the identification of bull and bear markets that is necessary for an evaluation of the real-time forecasts and we illustrate the aggregation performance of the various PCA techniques assuming knowledge of the full sample. To preserve space, we do not present any in-sample forecasts of stock market regimes (and returns). We also do not interpret the principal components, since the number of components and their interpretation might be different when considering a training set that only covers part of the full sample.

4.1. Classification of bull and bear markets

Despite its practical importance and relevance, there is no uniform definition of what exactly characterizes a bull or bear market (Gonzalez et al. 2006). In general, a stock market contraction is a persistent price decline associated with higher fluctuations. However, there is no consensus on how long such a period should last or how strong the price decline should be. We follow the literature (e.g., Kole and van Dijk 2017) and use dating rules for an evaluation of our real-time forecasts.

The underlying idea is to identify local peaks and troughs in the stock price series P_t of the S&P 500 without any distributional assumptions. The identified extreme points mark the turning points of the stock market and the period between a high (low) point and a low (high) point reflects a bear (bull) market. We follow the dating rule of Lunde and Timmermann (2004), as it focuses on absolute price changes and, thus, allows for an intuitive and transparent distinction. Their identification procedure (LT, henceforth) can be summarized as follows:

1. Given that the last observed extreme was a local maximum, referred to as P^{max}, the subsequent price series is checked against the following criteria:
   (a) The peak is updated if the stock market has risen above the last peak.
   (b) A local minimum has been found if the stock market has fallen by 10% or more.
   (c) There are no updates if neither a) nor b) took place.
2. Given that the last observed extreme was a local minimum, referred to as P^{min}, the subsequent price series is checked against the following criteria:
   (a) The trough is updated if the stock market has dropped below the last minimum.
   (b) A peak has been found if the stock market has risen by 15% or more.
   (c) There are no updates if neither a) nor b) took place.

In simple terms, periods that result in at least a 10% drop in stock prices are classified as bearish. A switch to a bull market follows if the stock price increase from the low is at least 15%. The particular thresholds are indeed arbitrary, but common in practice.

Fig. 2 and Table 2 show the log of the S&P 500 within bullish and bearish market regimes as identified by the LT filter. The biggest drop was caused by the GFC in 2007–2008 (−49%), whereas the bursting of the dotcom bubble (March 2000 to September 2001) marked the longest bear market with a duration of 78 weeks. The recent Covid-19 crash (February to March 2020) is historically the shortest contraction period, but the one with the third largest price slump.

During our evaluation period, the four economic recessions (according to the definition by the National Bureau of Economic Research, NBER) are always accompanied by a stock market contraction.[8] Despite the fact that the

[7] The starting point is restricted by data availability for many of the predictors, such as the forecasts from Consensus Economics, but also the VIX, corporate bond yields, credit spreads, and sentiment indicators.

[8] The first recession lasted from August 1990 to March 1991, the second from April to November 2001, the third from January 2008 to June 2009, and the most recent one (as of November 2021) from March to April 2020 (https://fred.stlouisfed.org/series/USRECD).

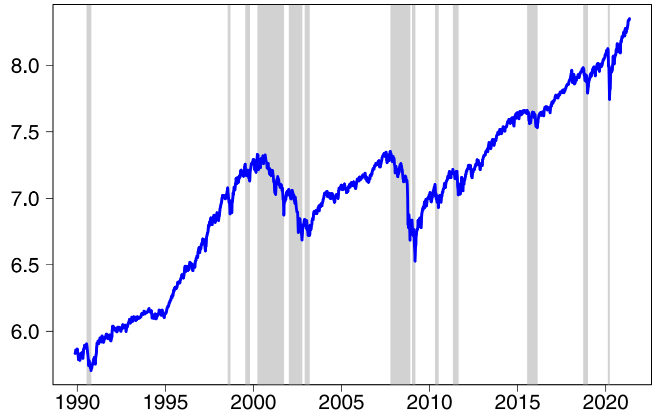

Fig. 2. Full-sample bull and bear market classification.
Notes: The figure shows the log S&P 500 price index and the identified bear markets as gray-shaded areas. The classification follows the dating rule of Lunde and Timmermann (2004)

Table 2
Bull and bear market periods.

Bull markets Dates	Durat.	Amplit.	Bear markets Dates	Durat.	Amplit.
1989-11-17 to 1990-07-13	35	8	1990-07-20 to 1990-10-12	13	−18
1990-10-19 to 1998-07-17	405	296	1998-07-24 to 1998-09-04	7	−18
1998-09-11 to 1999-07-16	45	46	1999-07-23 to 1999-10-15	13	−12
1999-10-22 to 2000-03-24	23	22	2000-03-31 to 2001-09-21	78	−37
2001-09-28 to 2002-01-04	15	21	2002-01-11 to 2002-10-04	39	−32
2002-10-11 to 2002-11-29	8	17	2002-12-06 to 2003-03-07	14	−11
2003-03-14 to 2007-10-12	240	88	2007-10-19 to 2008-11-21	58	−49
2008-11-28 to 2009-01-02	6	16	2009-01-09 to 2009-03-06	9	−27
2009-03-13 to 2010-04-23	59	78	2010-04-30 to 2010-07-02	10	−16
2010-07-09 to 2011-04-29	43	33	2011-05-06 to 2011-08-19	16	−18
2011-08-26 to 2015-07-17	204	89	2015-07-24 to 2016-02-12	30	−12
2016-02-19 to 2018-09-21	136	57	2018-09-28 to 2018-12-21	13	−18
2018-12-28 to 2020-02-14	60	40	2020-02-21 to 2020-03-20	5	−32
2020-03-27 to 2021-05-07	59	84			

Notes: The classification follows the dating rule of Lunde and Timmermann (2004). The duration is measured in weeks and the amplitude as the percentage price change between two subsequent extreme points.

duration and the amplitude of bear markets vary considerably, we can confirm that the stock market acts as an important leading indicator for the business cycle (Estrella and Mishkin 1998, Hamilton and Lin 1996). However, the stock market would predict even more recessions (see Chauvet and Potter 2000), displaying the "excess" sensitivity of expectations and risk aversion to bad news. Overall, the LT dating rule is able to detect persistent downward and upward trends as well as temporary bear market rallies (or short-run bull markets). Hence, it serves as a good proxy to evaluate the accuracy of the real-time predictions.

4.2. Data aggregation

To utilize the information from a high-dimensional dataset of potential predictors, we apply four different PCA techniques to aggregate the information into a few components and to filter out the noise: (i) conventional PCA, (ii) sparse PCA (SPCA), (iii) targeted PCA (TPCA), and (iv) targeted sparse PCA (TSPCA).

Table 3 shows the number of selected components and the proportion of explained variance under full-sample knowledge.

Six components are selected for the conventional PCA and the sparse PCA. These capture 56% (PCA) and 51% (SPCA) of the total variation. The benefits of soft thresholding become evident when targeting the (sparse) PCA on the VIX. In this case, three components are sufficient and explain more of the variation than the first three components of their non-targeted counterparts (TPCA: 46% vs. PCA: 42%; TSPCA: 41% vs. SPCA: 37%). When targeting on the equity risk premium (ERP), the explained variance of the first three components is similar to their non-targeted counterparts (TPCA: 41%; TSPCA: 38%).

Figures B1–B3 in Appendix B show the principal components over time. It is noticeable that the sparse PCA (right panel) achieves a more distinct smoothing over the indicators compared to the conventional PCA (left panel), irrespective of whether the set of predictors is unrestricted (Figure B1) or targeted (Figures B2 and B3). Hence, we can conclude that the sparse factors are more capable of filtering out the noise, confirming the results of Rapach and Zhou (2019).

As mentioned above, the number of obtained principal components (and their interpretation) might vary when considering a training set that only covers part of the full sample. Figure B4 in Appendix B provides an overview on

Table 3
Cumulative proportion of explained variance (in-sample).

	Full Dataset ($N = 146$)		VIX: Targeted dataset ($N = 75$)		ERP: Targeted dataset ($N = 75$)	
	PCA	Sparse PCA	PCA	Sparse PCA	PCA	Sparse PCA
PC1	0.19	0.15 (58)	0.25	0.20 (44)	0.19	0.18 (28)
PC2	0.35	0.30 (42)	0.37	0.31 (30)	0.34	0.32 (40)
PC3	0.42	0.37 (51)	0.46	0.41 (47)	0.41	0.38 (41)
PC4	0.47	0.42 (36)				
PC5	0.52	0.47 (40)				
PC6	0.56	0.51 (54)				

Notes: The number of principal components is based on the selection procedure presented in Section 2.1 and for the full sample period. For the sparse PCA, the proportion of explained variance is calculated via the QR decomposition of the (correlated) principal component scores. "Targeted dataset" refers to the subset of indicators obtained via soft thresholding. The number of non-zero coefficients of the sparse PCA is given in parentheses.

the number of principal components used in the out-of-sample exercise in Section 5. The number of components used in PCA and SPCA varies between five and eight, whereas for TPCA and TSPCA, three and seven mark the lower and upper bound.

5. Out-of-sample results

We use a recursive forecasting procedure with an expanding window to capture the stock market dynamics from October 22, 2004 to May 7, 2021, yielding a total of 864 forecasts. Our out-of-sample period starts with a prolonged bullish market (see Table 2). Starting from October 2007 onward, we have a total of 14 turning points that our models aim to predict in a real-time setting. Our entire methodology (estimation of PCAs, MS models, and forecast combination) is always applied on a weekly updated training sample. The first training set uses the available information from November 17, 1989 to October 15, 2004 to forecast regimes and returns for October 22, 2004. For the last forecast, information up to April 30, 2021 is used.

We evaluate the predictive power of our approach in terms of its statistical quality and its practical use for an investor. Our investment universe comprises a risky asset (SPDR S&P 500 ETF, Code: SPY) and an (almost) risk-free asset (3 M Treasury bill, secondary market rate). We resort to actually traded products to enable an assessment from an investor's perspective. For an evaluation of the economic value, we rely on two investment strategies. For regime forecasts, we employ a switching strategy, which allocates the total wealth either in the broad stock market or in Treasury bills, according to the different forecast clusters from Table 1. In the case of return forecasts, we utilize a mean–variance strategy where the stock market portfolio weight depends on the optimal conditional portfolio rule (Merton 1969). Short-selling and leverage are not allowed in both cases. As benchmarks, we utilize (simple) strategies based on the one-year moving average (MA_12M) for the regime forecasts. For the return forecasts, the historical average (HIST) of the equity risk premium serves as a benchmark. An MS model with TCTP according to Eq. (5) is applied for both types of forecasts. For an evaluation of the economic value, we additionally employ the straightforward buy-and-hold strategy (BH), the 50/50 strategy (50% equity and 50% risk-free), and the 60/40 strategy (60% equity and 40% risk-free) as further benchmarks.

Finally, we account for transaction costs to obtain a realistic perspective for an investor. Estimating transaction costs is not an easy task and depends on many factors (e.g., order size, market liquidity, and investor characteristics). We refrain from delving deeper into this topic and assume transaction costs of 20 basis points (bps) that are proportional to the size of the position change (e.g., Çakmaklı and van Dijk 2016) for the baseline case. The impact of alternative transaction cost assumptions (0 bps and 50 bps) as well as an ex ante consideration of transaction costs in spirit of Dal Pra et al. (2018) is investigated in Appendix E. All evaluation measures for the statistical quality and the economic performance are explained in Appendix C.

5.1. Regime predictability

Statistical performance: In the context of stock market regime identification, the timely detection of bear markets is particularly important for loss reduction. Put differently, the statistical evaluation follows the methodology of classification decisions. Hence, we have to handle the tradeoff between the true positive rate (i.e., a bear market is correctly predicted) and the false positive rate (i.e., a bull market is misclassified as a bear market; false alarm). Within a two-regime case, one typically relies on a cutoff of 50% in the predicted probabilities to differentiate between regimes. This threshold appears to be the most intuitive choice at first glance. The receiving operating characteristic (ROC) curve is a more nuanced approach to evaluating classifications, as it considers a grid of thresholds and displays the benefits (true positive rate) and costs (false positive rate) of a classification model in a two-dimensional figure (Fawcett 2006). A popular way to aggregate the performance of the ROC curve into a single value is to calculate the area under the curve (AUC). Since the ROC curve is plotted on a unit square, the AUC takes values between 0 and 1, where 1 (0.5) corresponds to a perfect (random) classification.

Table 4
Regime forecasts: Statistical performance.

Forecast	QPS	AUC	Accuracy	Bear	Bull
MA_12M (D_t^{MA})	0.310		0.836	0.475	0.906
TCTP	0.324	0.830	0.788	0.773	0.791
A-OBS-AVE	**0.250**	**0.853**	0.800	0.738	0.812
A-PC-AVE	**0.281**	**0.834**	0.795	0.766	0.801
A-SPC-AVE	**0.268**	**0.845**	0.793	0.745	0.802
A-TPC-AVE	**0.259**	**0.850**	0.808	0.766	0.816
A-TSPC-AVE	**0.240**	**0.861**	0.808	0.738	0.822
A-OBS-BMA	**0.290**	0.820	0.785	0.624	0.816
A-PC-BMA	**0.275**	0.813	0.803	0.681	0.827
A-SPC-BMA	**0.233**	**0.859**	0.818	0.688	0.844
A-TPC-BMA	**0.276**	**0.842**	0.796	0.695	0.816
A-TSPC-BMA	**0.277**	**0.836**	0.796	0.660	0.823
B-OBS-AVE	**0.245**	**0.853**	0.810	0.731	0.826
B-PC-AVE	**0.274**	**0.838**	0.802	0.773	0.808
B-SPC-AVE	**0.262**	**0.846**	0.799	0.752	0.808
B-TPC-AVE	**0.253**	**0.851**	0.808	0.752	0.819
B-TSPC-AVE*	**0.238**	**0.863**	0.808	0.731	0.823
B-OBS-BMA	**0.302**	0.803	0.774	0.582	0.812
B-PC-BMA	**0.285**	0.810	0.800	0.688	0.822
B-SPC-BMA	**0.256**	**0.846**	0.810	0.738	0.824
B-TPC-BMA	**0.279**	0.824	0.788	0.638	0.817
B-TSPC-BMA	**0.232**	0.824	0.819	0.638	0.855

Notes: All forecasts except D_t^{MA} (naïve 12-month moving average) and TCTP (MS model with TCTP and without external predictors) are estimated as MS-TVTP models. Specification A contains predictors in the switching equation and the conditional mean equation. Specification B contains predictors in the switching equation only. OBS: observable predictors; PC: conventional PCA; SPC: sparse PCA; TPC: conventional PCA with soft thresholding; TSPC: sparse PCA with soft thresholding; AVE: simple average; BMA: Bayesian model averaging; QPS: quadratic probability score; AUC: area under the curve; Accuracy: share of correctly predicted regimes overall (50% threshold); Bear/Bull: share of correctly predicted bearish/bullish regimes (50% threshold). See Appendix C for a detailed description of the evaluation measures. Forecasts that outperform both benchmarks are highlighted in bold. * indicates the best forecast according to an AUC test (DeLong et al., 1988).

As a naïve benchmark for regime identification, we consider the one-year unweighted moving average (MA) of the excess stock returns. The MA is often applied as an indicator to signal trends and is therefore useful for market timing decisions (see, among others, Brock et al. 1992). To separate the smoothed performance into two regimes, we define the binary variable D_t^{MA}:

$$D_t^{MA} = 0 \quad \text{if} \quad MA_t \geq 0 \quad \text{as bullish phase}$$
$$D_t^{MA} = 1 \quad \text{if} \quad MA_t < 0 \quad \text{as bearish phase}$$

The window length of one year is indeed arbitrary, but common in practice. A shorter length might lead to too many turning points and very short-lived bullish and bearish periods, whereas a longer memory would not appropriately account for the most recent price dynamics.[9]

Table 4 shows the statistical performance of our forecasts against the moving average and the simple MS model. All proposed models can outperform the MA and the MS model with TCTP in terms of the quadratic probability score (QPS). In addition, the AUC statistics are better for all forecasts (exceptions: A-OBS-BMA, A-PC-BMA, B-OBS-BMA, and B-PC-BMA) than for the MS model with TCTP. The total accuracy reaches up to 81.0% (B-OBS-AVE and B-SPC-BMA) with bear market accuracy rates of up to 77.3% (B-PC-AVE). However, our forecasts cannot consistently outperform both benchmarks in the classification metrics. Nevertheless, they perform better than the MA (the MS model with TCTP) when it comes to predicting bear (bull) markets.

For a formal identification of the best forecasts, we rely on an AUC test for the accuracy of regime predictions (DeLong et al. 1988). Table D1 in Appendix D displays the results. The best forecast according to this test is B-TSPC-AVE, which outperforms all but one forecast, followed by A-TSPC-AVE, which beats all but two forecasts. Hence, soft thresholding on a sparse PCA outperforms all other aggregation techniques (and the observable predictors). In general, the simple average performs better than the more complex BMA (exception: A-SPC-BMA). Finally, the richer Specification A (conditional transitions and mean) does not outperform the more parsimonious Specification B (conditional transitions only).

In addition to the average performance of the forecasts, it is of particular interest to see how timely recessions are detected. Figures D1–D5 in Appendix D show the bear market probabilities that the different forecasts predict.

[9] We also provide results with window lengths of 3, 6, 24, and 36 months as part of our robustness tests in Tables E1 and E2 of Appendix E. The 12-month MA performs best in the case of statistical performance (exception: share of correctly predicted bear markets). In terms of the economic value, longer window lengths provide, on average, higher returns, while shorter lengths yield better tail-risk measures and a lower standard deviation. Hence, we also rely on a 12-month window length in this case, as it provides a good compromise between returns and tail risk.

Table 5
Identification of turning points in bull and bear markets.

	Bull → Bear			Bear → Bull		
	Best	Worst	TCTP	Best	Worst	TCTP
Global Financial Crisis I (2007-10-19 to 2008-11-21)	0	0	0	0	+6	+6
Global Financial Crisis II (2009-01-09 to 2009-03-06)	0	+1	0	0	+40	+41
Flash Crash Aftermath (2010-04-30 to 2010-07-02)	+1	+2	+1	+4	+13	+14
Debt Crisis (2011-05-06 to 2011-08-19)	+8	+12	+8	+3	+22	+23
Chinese Market Crash (2015-07-24 to 2016-02-12)	+4	+4	+4	0	+5	+5
Economic Slowdown Fear (2018-09-28 to 2018-12-21)	+2	+4	+2	+1	+7	+7
Covid-19 Crash (2020-02-21 to 2020-03-20)	+1	+1	+1	+3	+36	+18

Notes: The table shows the out-of-sample delay (in weeks) when identifying regime switches from bull to bear markets and from bear to bull markets. The dating rule of Lunde and Timmermann (2004) (assuming full-sample knowledge) is used for the classification of bull and bear markets. Across all forecast combinations, the performance of the best model and the worst model is reported with the delay of the MS-TCTP model as the benchmark. The threshold for the bear market probability is 50%.

Overall, the predicted bear market probabilities respond promptly to regime turning points. In addition, the respective regime forecasts have a high degree of similarity across the different forecast combination clusters. Again, the simple average outperforms the BMA in terms of the R^2 for all forecast combinations except those using a sparse PCA. In terms of this metric, forecasts A-SPC-BMA ($R^2 = 0.287$) and B-TSPC-AVE ($R^2 = 0.277$) perform slightly better than the remaining forecasts.

Table 5 aggregates the information from Figures D1–D5 and shows how quickly turning points are detected over the course of the different recessions. The MS model with TCTP and without external predictors serves as the benchmark. As an illustration, the best model can identify the start and end of the GFC without a delay. The Covid-19 crash is also classified as a bear market from the end of February 2020 onwards, with the re-entry taking place in mid-April. Confirming the impression from Table 4, the TCTP specification is well suited to detect bear markets. Even the best of our forecasts is only able to match its performance (see column "Bull → Bear"). However, a key advantage of our approach is to identify the turning point from bear to bull markets in a timely manner, as our best model never exceeds a delay of four weeks when classifying the switch into a bull market (see column "Bear → Bull").

Economic value: All the metrics so far have tested for (sometimes nuanced) differences in the statistical performance of the different forecasts. For an investor, however, it is important to see whether these statistical differences turn into economic value added, in particular when considering transaction costs. Hence, we evaluate the profitability of regime forecasts by translating the regime probabilities into a binary investment strategy that allocates the total wealth either to the stock market (risk-on) or to short-term government bonds (risk-off). If a bear (bull) market is predicted, we avoid (going long in) the stock market. Accordingly, the optimal stock market weight w_t^* goes hand in hand with a threshold dependent indicator $\hat{I}_t(\tau)$:

$$w_t^* = \hat{I}_t(\tau) \tag{11}$$

In our baseline scenario, we consider transaction costs ex post when calculating the performance metrics, and switch the indicator to zero if the bear market probability $\hat{p}_t$ exceeds a certain threshold τ (in our baseline scenario, we assume $\tau = 0.5$)[10]:

$$\hat{I}_t = 0 \quad \text{if} \quad \hat{p}_t \geq \tau$$
$$\hat{I}_t = 1 \quad \text{if} \quad \hat{p}_t < \tau$$

Table 6 shows the economic value of the forecasts in the baseline scenario against five benchmarks (BH, 50/50,

[10] As part of our robustness tests, we also consider a more (less) recession-averse agent and set $\tau = 0.25$ ($\tau = 0.75$). In our second scenario, we consider transaction costs ex ante in spirit of Dal Pra et al. (2018) in combination with varying threshold levels:

$$\hat{I}_t = 0 \quad \text{if} \quad \hat{p}_t \geq \tau \quad \textit{and} \quad -\hat{r}_{e,t+1} \geq c_t$$
$$\hat{I}_t = 1 \quad \text{if} \quad \hat{p}_t < \tau \quad \textit{and} \quad \hat{r}_{e,t+1} \geq c_t$$
$$\hat{I}_t = I_{t-1} \quad \textit{else}$$

Hence, an investor switches from risk-off (risk-on) to risk-on (risk-off) only if the predicted stock return (loss) that results from the regime forecast exceeds the transaction costs c_t. In all other cases, no trade is executed. The economic performance for an ex ante consideration of different transaction costs is documented in Table E3 in Appendix E. In Table E4, we vary the amount of ex post transaction costs and the recession thresholds. To conserve space, we show only the results for the annualized certainty equivalent return as an important risk-adjusted measure.

We can infer a couple of interesting results from this robustness test. First, more trades are executed in the absence of transaction costs, which also leads to higher annualized certainty equivalent returns. Conversely, higher transaction costs lead to fewer trades and a lower CER. Second, when considering transactions costs ex ante, the CER is, on average, lower as compared to the case of ex post trading costs, as fewer trades are conducted. 50 bps appears to be prohibitive for the case of ex ante transaction costs. Third, lowering the threshold makes the positioning "too cautious" and leads to a disproportionate decline in returns relative to the risk improvement. A cutoff point of 75%, however, leads to higher annualized CER returns.

Table 6
Regime forecast: Economic value.

Forecast	R^{cum}	$\bar{R}$	$\bar{\sigma}$	SR	CER	MaxDD	VaR	CVaR
BH	5.28	10.53	17.90	0.51	4.38	−54.6	−3.87	−6.04
50/50	2.76	6.30	8.57	0.58	3.90	−29.2	−1.95	−2.90
60/40	3.18	7.22	10.35	0.57	4.31	−34.7	−2.34	−3.51
MA_12M (D_t^{MA})	4.03	8.75	12.92	0.57	4.92	−22.2	−2.79	−4.58
TCTP	3.96	8.63	10.25	0.71	5.74	−12.5	−2.11	−3.55
A-OBS-AVE	3.87	8.49	10.43	0.69	5.53	−12.5	−2.16	−3.66
A-PC-AVE	3.77	8.32	10.30	0.68	5.41	−17.7	−2.12	−3.59
A-SPC-AVE	4.16	8.95	10.40	**0.73**	**6.01**	−12.5	−2.14	−3.61
A-TPC-AVE	4.21	9.03	10.58	**0.73**	**6.03**	−17.7	−2.14	−3.65
A-TSPC-AVE	4.18	8.99	10.57	**0.72**	**5.99**	−14.4	−2.16	−3.64
A-OBS-BMA	2.48	5.61	10.72	0.40	2.61	−23.7	−2.39	−3.93
A-PC-BMA	4.70	9.76	10.92	**0.77**	**6.64**	−17.5	−2.16	−3.73
A-SPC-BMA	4.29	9.16	11.39	0.69	**5.89**	−26.5	−2.39	−3.97
A-TPC-BMA	3.89	8.52	11.14	0.65	5.34	−17.5	−2.31	−3.88
A-TSPC-BMA	4.08	8.84	10.78	0.70	**5.77**	−18.9	−2.22	−3.71
B-OBS-AVE	4.24	9.08	10.55	**0.74**	**6.08**	−12.5	−2.14	−3.65
B-PC-AVE	3.99	8.69	10.22	**0.72**	**5.80**	−12.5	−2.12	−3.50
B-SPC-AVE	4.37	9.29	10.49	**0.76**	**6.31**	−12.5	−2.14	−3.61
B-TPC-AVE	4.30	9.18	10.60	**0.74**	**6.17**	−15.3	−2.14	−3.64
B-TSPC-AVE	4.45	9.40	10.61	**0.76**	**6.38**	−13.2	−2.16	−3.64
B-OBS-BMA	2.64	6.01	10.64	0.44	3.02	−18.0	−2.30	−3.86
B-PC-BMA	4.75	9.84	10.70	**0.80**	**6.79**	−13.0	−2.14	−3.63
B-SPC-BMA*	5.06	10.25	10.86	**0.82**	**7.14**	−12.6	−2.16	−3.62
B-TPC-BMA	3.32	7.49	11.01	0.56	4.36	−17.5	−2.41	−3.94
B-TSPC-BMA	4.82	9.92	11.39	**0.75**	**6.64**	−24.3	−2.33	−3.83

Notes: The table shows the economic value of different investment strategies. BH: Buy & Hold Strategy of the S&P 500; 50/50 and 60/40: mixed strategy S&P 500 and 3 M Treasury bill; D_t^{MA}: naïve 12-month moving average; TCTP: MS model with TCTP and without external predictors. Specification A contains predictors in the switching equation and the conditional mean equation. Specification B contains predictors in the switching equation only. OBS: observable predictors; PC: conventional PCA; SPC: sparse PCA; TPC: conventional PCA with soft thresholding; TSPC: sparse PCA with soft thresholding; AVE: simple average; BMA: Bayesian model averaging; R^{cum}: final wealth of strategy assuming a $1 investment; $\bar{R}$: annualized average returns; $\bar{\sigma}$: annualized standard deviation; SR: annualized Sharpe ratio; CER: annualized certainty equivalent return with $\gamma = 3$; MaxDD: maximum drawdown; VaR: value at risk; CVaR: conditional value at risk. See Appendix C for a detailed description of the evaluation measures. Forecasts that outperform all benchmarks are highlighted in bold. * indicates the best forecast according to the risk-adjusted performances measures (SR and CER).

60/40, MA, and TCTP). The buy-and-hold strategy performs best in terms of final wealth (R^{cum}) and annualized average returns ($\bar{R}$). The 50/50 strategy yields the lowest annualized standard deviation ($\bar{\sigma}$) and the best (conditional) value at risk (VaR and CVaR). Our forecasts, in turn, perform particularly well when considering the annualized Sharpe ratio (SR) and the annualized certainty equivalent return (CER). In particular, eight forecasts that rely on the restricted Specification B perform better than all benchmarks in these two risk-adjusted measures. Hence, together with the similar statistical performance of Specification A and Specification B, these results indicate that modeling the conditional transitions might be sufficient when forecasting regimes.

Turning to the forecast combination scheme, we find that more combinations based on the simple average—as opposed to the BMA—outperform the five benchmarks. On the other hand, the best risk-adjusted metrics are found for the forecast combination B-SPC-BMA. The latter finding is reassuring, since B-SPC-BMA is also the forecast with the best statistical accuracy. In addition, the models with the best AUC from Table 4 (A-TSPC-AVE and B-TSPC-AVE) outperform all benchmarks in terms of their SR and CER. Hence, their statistical accuracy is also reflected in actual value added.

Another takeaway is the worse performance of observable predictors in comparison to the principal components, confirming the findings of Neely et al. (2014) and Çakmaklı and van Dijk (2016). Within PCA models, introducing sparsity in PCA helps to improve the predictability with, on average, a better risk-adjusted performance. In addition to a more straightforward interpretation, the sparse factors provide a sharper distinction between signal and noise, which is in line with the results of Rapach and Zhou (2019). Finally, it is worth noting that our best forecast (B-SPC-BMA) in terms of risk-adjusted performance comes close to the return metrics of the buy-and-hold strategy, while at the same time having a much lower annualized standard deviation ($\bar{\sigma}$) and a maximum drawdown that is close to the best benchmark value (MS model with TCTP).

Another way of illustrating the economic performance of the different forecasts is to plot the cumulative returns of the different strategies over time. This provides another view of the ability to detect turning points in a timely manner. Figure D6 in Appendix D shows the cumulative returns over time. As indicated by the results in Table 6, the BH strategy performs best when considering the final wealth. However, in particular during the GFC and the Covid-19 crash, losses can be reduced and re-entry points can be found in a timely manner when relying on our forecast methodology. Finally, it also becomes evident that our best forecast (B-SPC-BMA) performs even better than the BH strategy for almost the entire out-of-sample period and is only outperformed during the booming post-Covid-crash stock market, which shifts the overall results in favor of the BH strategy. All these findings suggest that it pays off to model the switching process with TVTP when evaluating returns on a risk-adjusted basis.

5.2. Return predictability

Statistical performance: When it comes to forecasting stock market returns, accurate point forecasts in terms of a low mean squared prediction error (MSPE) are difficult to find, in particular at a weekly frequency. It is therefore common to compare the forecast quality relative to the historical average. Hence, we rely on the R^2_{OS} proposed by Campbell and Thompson (2008). The historical average is calculated with an expanding window, so that the period from November 17, 1989 to October 15, 2004 is used for the first forecast.[11]

Table 7 shows the statistical performance of the return forecasts. There is no added value with regard to the R^2_{OS} and RMSE over the entire out-of-sample period. The values for R^2_{OS} are negative (except for A-TPC-AVE and A-TPC-BMA) and the null hypothesis of $R^2_{OS} \leq 0$ cannot be rejected. The sign predictability varies between 55% and 58%; positive returns are correctly predicted in 76% to 83% of the cases, and the negative return accuracy rate ranges from 21% to 30% (when excluding the outlier forecast B-OBS-BMA that predicts 94% of the positive returns and 7% of the negative returns).

For a formal identification of the best forecasts, we rely on the (Diebold & Mariano, 1995) test for the accuracy of return predictions. Table D2 in Appendix D displays the results. A-TPC-AVE is the best forecast, as it outperforms nine other forecasts, followed by B-TSPC-BMA (8/20). The best forecast for regime predictions (B-SPC-BMA) ranks third (6/20). According to this test, Specification B, which models only the conditional transitions, performs slightly better than Specification A, which also models the conditional mean process. Applying sparsity and/or soft thresholding on the PCA appears to outperform the observable predictors and a conventional PCA. Finally, considering the forecast combination scheme, Specification A (B) performs worse when relying on the BMA (simple average).[12]

Economic value: As a final step, we test whether the dispiriting statistical performance of the return forecasts is also reflected in their economic value. We assume a risk-averse agent with mean–variance preferences. Solving the standard expected utility maximization, we obtain the following optimal stock market weight that is restricted between 0% and 100%:

$$w_t^* = \frac{1}{\gamma}\frac{\hat{r}_{t+1}}{\hat{\sigma}^2_{t+1}} \qquad (12)$$

$\hat{r}_{t+1}$ represents the one-step-ahead return forecast and $\hat{\sigma}^2_{t+1}$ the expected variance. The coefficient of relative risk aversion γ is set to 3, and we use the historical five-year variance as a risk proxy.[13] We consider proportional ex post transaction costs of 20 bps when calculating the performance metrics.[14]

Table 8 shows the economic value of the forecasts against five benchmarks (BH, 50/50, 60/40, HIST, and TCTP). Again, the buy-and-hold strategy performs best in terms of final wealth (R^{cum}) and annualized average returns ($\bar{R}$). The 50/50 strategy yields the lowest annualized standard deviation ($\bar{\sigma}$) and the best conditional value at risk (CVaR). A few of our forecasts outperform all benchmarks in terms of the risk-adjusted metrics (SR and CER) and tail-risk measures (MaxDD and VaR). These all rely on the restricted Specification B. Hence, it appears that modeling conditional returns does not improve the economic value of regime or return forecasts.

Turning to the combination schemes, we find that forecast combinations using the BMA (again including B-SPC-BMA) produce slightly better values for the risk-adjusted performance measures than the simple average. Again, the aggregation of information in principal components is more helpful than just relying on observable predictors. Finally, and perhaps the most important takeaway, the return forecasts provide a lower final wealth, lower annualized returns, a lower SR, and a lower CER when compared to the regime forecasts (see Table 6). On

11 The assumption of an expanding window is required for an evaluation of our forecasts within a nested framework and for the same training sample. Obviously, such a benchmark model implies that returns (and regimes) are unpredictable. However, when comparing the RMSE of the historical average (2.49) to moving averages with varying lengths, it is never outperformed (3 M: 2.59; 6 M: 2.53; 12 M: 2.51; 24 M: 2.50; 36 M: 2.49).

12 To conserve space, we do not show a graphical representation of the excess return forecasts over time. These are very noisy and not informative, because their R^2 never exceeds 0.005. All omitted results are available on request.

13 It has to be noted that the forecast results are also influenced by the expected variance proxy. A rolling window of five years implies a high degree of persistence. Consequently, the stock exposure might be biased downwards, in particular after large shocks. An extension for future research would be to explicitly include a variance forecast.

14 Although the assumption of $\gamma = 3$ is common in the empirical literature, we provide results for $\gamma = 2$ and $\gamma = 5$ in Table E5 of Appendix E. In addition, we vary the amount of ex post transaction costs. Including ex ante transaction costs in a mean–variance optimization is a more complex task that we leave open for future studies. Again, we show only the results for the annualized certainty equivalent return as an important risk-adjusted measure. We can infer a couple of interesting results from this robustness test. First, more trades are executed in the absence of transaction costs, which also leads to higher annualized certainty equivalent returns. Yet, we do not find a substantial improvement of our forecasts in terms of outperforming the benchmarks. Second, increasing the degree of risk aversion (to 5) makes the positioning "too cautious" and leads to a disproportionate decline in returns relative to the risk improvement.

Finally, we also use our excess return prediction as a trigger for the switching strategy. Following Dal Pra et al. (2018), we fully invest in the stock market whenever the excess returns are positive or at least zero. For negative predictions (i.e., expected stock returns are smaller than the risk-free rate), we only invest in Treasury bills. Table E6 in Appendix E shows the results. Here, we find an outperformance of the benchmarks in the absence of transaction costs or when transaction costs are only considered ex post. In particular, Specification B, which models only the conditional transitions, performs well in that regard.

Table 7
Return forecasts: Statistical performance.

Forecast	RMSE	CW		Direction	R^+	R^-
		R^2_{OS}	p-val.			
HIST	2.4866			0.582	1.000	0.000
TCTP	2.4929	−0.511	0.60	0.549	0.740	0.283
A-OBS-AVE	2.4918	−0.418	0.47	0.567	0.789	0.258
A-PC-AVE	2.4931	−0.527	0.53	0.553	0.779	0.238
A-SPC-AVE	2.4940	−0.596	0.58	0.552	0.783	0.230
A-TPC-AVE*	**2.4854**	**0.096**	0.18	0.576	0.825	0.230
A-TSPC-AVE	2.4952	−0.695	0.71	0.571	0.831	0.208
A-OBS-BMA	2.5120	−2.058	0.90	0.558	0.807	0.211
A-PC-BMA	2.5035	−1.371	0.60	0.565	0.807	0.227
A-SPC-BMA	2.5132	−2.152	0.87	0.551	0.763	0.255
A-TPC-BMA	**2.4826**	**0.318**	0.20	0.566	0.759	**0.296**
A-TSPC-BMA	2.5038	−1.388	0.72	0.571	0.807	0.241
B-OBS-AVE	2.4917	−0.412	0.61	0.566	0.797	0.244
B-PC-AVE	2.4941	−0.611	0.65	0.552	0.759	0.263
B-SPC-AVE	2.4926	−0.483	0.58	0.556	0.758	0.274
B-TPC-AVE	2.4924	−0.474	0.57	0.560	0.777	0.258
B-TSPC-AVE	2.4926	−0.486	0.61	0.560	0.781	0.252
B-OBS-BMA	2.4906	−0.322	0.88	0.579	0.944	0.069
B-PC-BMA	2.4949	−0.676	0.66	0.559	0.775	0.258
B-SPC-BMA	2.4900	−0.275	0.47	0.559	0.767	0.269
B-TPC-BMA	2.4949	−0.675	0.70	0.565	0.779	0.266
B-TSPC-BMA	2.4896	−0.243	0.48	0.565	0.817	0.213

Notes: HIST: historical average of excess stock returns; TCTP: MS model with TCTP and without external predictors. Specification A contains predictors in the switching equation and the conditional mean equation. Specification B contains predictors in the switching equation only. OBS: observable predictors; PC: conventional PCA; SPC: sparse PCA; TPC: conventional PCA with soft thresholding; TSPC: sparse PCA with soft thresholding; AVE: simple average; BMA: Bayesian moving average; R^2_{OS}: out-of sample R^2; CW: CW test statistic; Direction: correctly predicted forecast direction; R^+: true positive forecasts; R^-: true negative forecasts. See Appendix C for a detailed description of the evaluation measures. Forecasts that outperform both benchmarks are highlighted in bold. * indicates the best forecast according to the test by Diebold and Mariano (1995).

the other hand, return forecasts based on the restricted Specification B provide a lower standard deviation and better tail-risk measures than the corresponding regime forecasts. Hence, there is some economic value added in forecasting returns, in particular for a risk-averse investor.

Next, we graphically inspect the ability of the forecasts to detect turning points in a timely manner. Figure D7 in Appendix D shows the cumulative returns over time. As mentioned above, the BH strategy performs best when considering the final wealth. However, in particular during the GFC and the Covid-19 crash, losses can be reduced and re-entry points can be found in a timely manner when relying on our forecast methodology and the restricted Specification B. But it does become evident once more that the cumulative performance of our forecasts is worse when predicting returns in comparison to predicting regimes (see Figure D6).

In a final step, we explore the performance of our forecast methodology conditional on the state of the stock market or the state of the business cycle. For this purpose, we separate the forecasts into two subsamples depending on whether the observation in $t+1$ is assigned to a bear market (recession) or a bull market (expansion). To classify the market state, we use the dating rule of Lunde and Timmermann (2004) (see Section 4.1) and for the business cycle, we rely on the binary recession indicator of the NBER. Table 9 shows the results where the historical mean is used as a benchmark for calculating the R^2_{OS} and the Δ_{CER}.

Confirming the broad consensus in the literature, return predictability is especially prevalent in "bad times". Here, the forecast-based strategy generates statistical and, in particular, economic value. A risk-averse mean–variance investor is willing to pay an annualized management fee (as indicated by Δ_{CER}) of up to 29.40% (in bear markets) and 39.10% (in recessions) to participate in the forecast-based strategy (B-TPC-BMA). The economic performance of B-SPC-BMA ranks second in that regard. In addition, the statistical performance is better during bear markets with significant values for the R^2_{OS} of more than 2% (mostly for Specification B). During economic recessions, we find positive but insignificant values for the R^2_{OS} (again, mostly for Specification B). However, we have to conclude that all forecast combinations are clearly inferior to the historical average during bull markets or expansions. All these results are in line with previous findings for regime-switching models (Henkel et al. 2011) and for return predictions in general (Rapach and Zhou 2013).

5.3. Discussion

Before concluding, we need to revisit our testing framework and, in particular, shed light on the difference in performance between regime forecasts and return forecasts. It is not possible to directly compare the statistical goodness of the two forecasts, due to the different scaling of the target variables (binary classification versus continuous scale). However, it is noticeable that the regime

Table 8
Return forecast: Economic value.

Forecast	R^{cum}	$\bar{R}$	$\bar{\sigma}$	SR	CER	MaxDD	VaR	CVaR
BH	5.28	10.53	17.90	0.51	4.38	−54.6	−3.87	−6.04
50/50	2.76	6.30	8.57	0.58	3.90	−29.2	−1.95	−2.90
60/40	3.18	7.22	10.35	0.57	4.31	−34.7	−2.34	−3.51
HIST	2.29	5.11	12.86	0.30	1.35	−43.5	−2.51	−4.77
TCTP	3.09	7.03	9.46	0.61	4.39	−12.9	−1.97	−3.30
A-OBS-AVE	2.21	4.88	11.39	0.32	1.66	−34.0	−2.07	−4.06
A-PC-AVE	2.81	6.42	9.91	0.52	3.65	−19.7	−2.11	−3.55
A-SPC-AVE	2.51	5.70	10.14	0.44	2.88	−28.0	−2.10	−3.70
A-TPC-AVE	2.27	5.06	12.32	0.31	1.51	−37.6	−2.16	−4.45
A-TSPC-AVE	1.86	3.81	11.82	0.22	0.45	−42.0	−2.12	−4.40
A-OBS-BMA	1.69	3.21	13.57	0.14	−0.81	−43.6	−2.12	−4.88
A-PC-BMA	2.30	5.13	12.08	0.32	1.66	−39.7	−2.32	−4.37
A-SPC-BMA	1.52	2.57	12.61	0.10	−1.06	−43.8	−2.23	−4.72
A-TPC-BMA	1.96	4.13	13.03	0.22	0.31	−36.5	−2.02	−4.53
A-TSPC-BMA	2.06	4.44	11.76	0.27	1.10	−40.5	−2.34	−4.39
B-OBS-AVE	3.04	6.93	9.26	0.61	4.34	**−12.4**	**−1.91**	−3.25
B-PC-AVE	3.03	6.90	9.39	0.60	4.28	−13.6	−1.95	−3.29
B-SPC-AVE	3.02	6.87	9.35	0.60	4.26	−13.4	**−1.91**	−3.30
B-TPC-AVE	3.09	7.03	9.43	0.61	**4.40**	−13.9	**−1.91**	−3.31
B-TSPC-AVE	2.93	6.68	9.40	0.57	4.06	−15.1	−1.95	−3.34
B-OBS-BMA	2.02	4.31	10.22	0.30	1.47	−30.4	−2.22	−3.89
B-PC-BMA*	3.38	7.61	9.66	**0.65**	**4.90**	**−12.4**	**−1.91**	−3.43
B-SPC-BMA	3.08	7.01	9.28	**0.62**	**4.42**	−16.2	**−1.90**	−3.29
B-TPC-BMA	2.86	6.54	9.28	0.56	3.95	−13.5	**−1.87**	−3.32
B-TSPC-BMA	2.81	6.42	9.63	0.53	3.73	−25.7	−2.00	−3.44

Notes: The table shows the economic value of different investment strategies. BH: Buy & Hold Strategy of the S&P 500; 50/50 and 60/40: mixed strategy S&P 500/3 M Treasury bill; HIST: historical average of excess stock returns; TCTP: MS model with TCTP and without external predictors. Specification A contains predictors in the switching equation and the conditional mean equation. Specification B contains predictors in the switching equation only. OBS: observable predictors; PC: conventional PCA; SPC: sparse PCA; TPC: conventional PCA with soft thresholding; TSPC: sparse PCA with soft thresholding; AVE: simple average; BMA: Bayesian model averaging; R^{cum}: final wealth of strategy assuming a $1 investment; $\bar{R}$: annualized average returns; $\bar{\sigma}$: annualized standard deviation; SR: annualized Sharpe ratio; CER: annualized certainty equivalent return with $\gamma = 3$; MaxDD: maximum drawdown; VaR: value at risk; CVaR: conditional value at risk. See Appendix C for a detailed description of the evaluation measures. Forecasts that outperform all benchmarks are highlighted in bold. * indicates the best forecast according to the risk-adjusted performances measures (SR and CER).

forecasts outperform their benchmarks more often than the return forecasts do. A better comparability can be achieved through utility or profit analysis.

From an investor's point of view, it is striking that the strategy based on regime probabilities has a more attractive risk-return structure. Obviously, predicting trends is much easier than generating point forecasts. However, since regime probabilities are also a significant factor in the return forecast setting due to the weighting of the regime-dependent averages in Eq. (8), the difference in the performance cannot be entirely caused by this. Another reason might be excessive trading and large transaction costs. Table D3 in Appendix D shows the average annual turnover and the cumulative transaction costs of the two forecast strategies. These differ significantly only for Specification A. For Specification B, the cumulative transaction costs are similar. But even if we set transaction costs to zero, significant differences in the CERs remain (see column TC = 0 bps of Tables E4 and E5 in Appendix E). Hence, transaction costs cannot fully explain the differences in the economic performance of regime forecasts and return forecasts.

The profile of the respective investment strategy might be another key factor. To evaluate regime forecasts, a switching strategy is applied, which invests either only in Treasury bills or only in the stock market. Such a strategy (e.g., Dal Pra et al. 2018, Pesaran and Timmermann 1995) tends to reflect the investment behavior of a risk-neutral investor. In contrast, the mean–variance strategy explicitly takes risk aversion into account. Since there is no direct link between the regime probability and the classical utility function of an investor, we can only implicitly account for this fact by reducing the threshold to classify a bear market (or by reducing the relative risk aversion of the mean–variance strategy). Our robustness checks (see Tables E4 and E5 in Appendix E) show that lowering the threshold or the relative risk aversion narrows the gap between the performance of regime forecasts and return forecasts.

Another source affecting the degree of the stock allocation is the variance proxy. We rely on a common benchmark (sample variance of a rolling 5-year window) since we are not aiming at volatility forecasting. The resulting high variance persistence might bias the stock exposure downwards for the mean–variance strategy and the return forecasts, in particular after significant shocks. Conversely, when calculating the regime probabilities to obtain inferences about the current state, an estimate of the regime-dependent variance is used to determine the density functions. Since the second moment is crucial for identifying regime shifts (Kole and van Dijk 2017), a certain proportion of the difference in performance could be

Table 9
Market and economic state dependency of return forecasts.

	LT dating rule				NBER			
	Bull		Bear		Expansion		Recession	
	R^2_{OS}	Δ_{CER}	R^2_{OS}	Δ_{CER}	R^2_{OS}	Δ_{CER}	R^2_{OS}	Δ_{CER}
TCTP	−2.77	−4.52	**2.06**	27.86	−1.22	−1.44	0.42	36.17
A-OBS-AVE	−3.29	−5.47	**2.86**	17.94	−0.44	−1.93	−0.38	15.60
A-PC-AVE	−3.15	−4.45	**2.47**	24.16	−0.91	−1.32	−0.03	28.68
A-SPC-AVE	−3.05	−4.66	**2.20**	21.27	−0.85	−1.48	−0.26	23.22
A-TPC-AVE	−2.04	−4.16	**2.53**	12.59	−0.42	−0.90	0.78	7.21
A-TSPC-AVE	−2.24	−4.33	**1.07**	9.22	−0.53	−1.08	−0.92	1.34
A-OBS-BMA	**0.60**	−2.60	−5.09	−1.19	0.00	−2.30	−4.78	−3.13
A-PC-BMA	−4.21	−3.20	**1.86**	10.63	−1.51	−0.76	−1.19	7.81
A-SPC-BMA	−3.15	−5.38	−1.01	6.32	−1.23	−3.13	−3.37	2.06
A-TPC-BMA	−0.91	−4.62	1.71	9.26	−0.76	−3.47	1.74	14.48
A-TSPC-BMA	−2.80	−4.05	0.22	11.45	−1.25	−1.80	−1.57	10.25
B-OBS-AVE	−2.63	−4.56	**2.12**	27.77	−0.96	−1.31	0.31	34.67
B-PC-AVE	−3.05	−4.57	**2.17**	27.54	−1.28	−1.53	0.27	35.91
B-SPC-AVE	−3.00	−4.68	**2.39**	27.85	−1.21	−1.61	0.48	36.35
B-TPC-AVE	−2.94	−4.32	**2.33**	27.19	−1.14	−1.32	0.40	35.24
B-TSPC-AVE	−2.87	−4.45	**2.23**	26.07	−1.04	−1.33	0.25	32.28
B-OBS-BMA	−1.34	−4.95	**0.84**	15.84	−0.52	−2.52	−0.06	19.03
B-PC-BMA	−3.24	−3.88	**2.25**	27.76	−1.42	−1.05	0.31	37.65
B-SPC-BMA	−2.77	−4.69	**2.57**	28.66	−1.07	−1.72	0.78	38.76
B-TPC-BMA	−3.11	−5.51	**2.11**	29.40	−1.26	−2.29	0.10	39.10
B-TSPC-BMA	−2.21	−4.20	**2.00**	23.69	−0.84	−1.32	0.54	29.33

Notes: The table shows the statistical performance and the economic value of the return forecasts in different states of the stock market and the economy (all in %). Market states are classified with the LT dating rule described in Section 4.1. Economic states are defined by the NBER classification. R^2_{OS}: out-of-sample R^2; Δ_{CER}: difference in the annualized certainty equivalent return (with $\gamma = 3$) between the forecast-based strategy and the historical average. TCTP: MS model with TCTP and without external predictors. Specification A contains predictors in the switching equation and the conditional mean equation. Specification B contains predictors in the switching equation only. OBS: observable predictors; PC: conventional PCA; SPC: sparse PCA; TPC: conventional PCA with soft thresholding; TSPC: sparse PCA with soft thresholding; AVE: simple average; BMA: Bayesian model averaging. See Appendix C for a detailed description of the evaluation measures. Forecasts that outperform the historical average according to the (Clark & West, 2007) statistic with a 10% significance level are highlighted in bold.

attributed to this procedure. However, if the explicit consideration of risk aversion as well as the proxy of the variance were to explain the different performance of regime and return forecast strategies, a switching strategy based on return forecasts should show very similar performance to its counterpart based on regime forecasts. Table E6 in Appendix E shows the result of a switching strategy (Dal Pra et al. 2018) that invests in stocks (Treasury bills) whenever the prediction is positive or zero (negative). We find that the CERs are decreasing in the vast majority of cases when compared to Tables E3 and E4.

Even if the difference in performance can be partly explained by the considerations above, regime forecasts remain superior to return forecasts in terms of their economic value. They are able to capture the trend-changing behavior of markets so that tail risks are reduced without sacrificing large returns. However, it is worth noting that the return predictions add statistical and, in particular, economic value during recessions or in declining markets.

6. Conclusions

Using a high-dimensional dataset of macro-financial variables, this paper offers a promising approach to predicting stock market regimes on a weekly basis. Since stock market predictions suffer particularly from parameter instability and model uncertainty, our approach combines the merits of dimensionality reduction techniques, regime-switching models, and forecast combination. We provided a comprehensive overview of the empirical usefulness of Markov-switching models with principal components and time-varying transition probabilities.

Our best weekly regime forecasts use a (targeted) sparse principal component Markov-switching model and time-varying transition probabilities. They are suitable to respond to trend changes in a timely manner, either to participate in recoveries or to prevent losses. This is also reflected in actual economic value added, as many of our forecasts exceed all benchmarks in risk-adjusted performance measures. However, when considering stock market returns, our forecasts do not statistically outperform common benchmarks. The fact that return forecasts perform worse than regime forecasts is not surprising. Predicting the broader trend of the stock market is obviously easier than providing point forecasts, in particular on a weekly basis. This outperformance can also be explained—to some extent—by differences in the testing procedure and the investment strategy. Nevertheless, our return forecasts still provide some economic value added for risk-averse investors, as they generate a lower annualized standard deviation of the returns and better tail-risk measures than the corresponding regime forecasts. We also confirmed previous findings that return predictability is limited to recessions or to periods of market turmoil.

In addition, we found that it is sufficient to model the time-varying conditional transitions in a Markov-switching model. Additionally modeling the conditional mean introduces further noise into the forecasts and particularly harms the economic performance of our forecasts. Based on our results, we suggest relying on dimensionality reduction techniques (instead of just relying on observable predictors) and enhancing the conventional principal component analysis with shrinkage methods such as sparsity and/or soft thresholding. Concerning the forecast combination technique, we did not find a clear advantage of Bayesian model averaging over the simple average.

Our results offer a variety of starting points for future work. First, modeling intra-regime dynamics in greater detail in the context of our forecasting model could be a promising extension—if practically feasible. Thereby, incorporating more than two regimes directly (Maheu et al. 2012) or sequential partitions (Hauptmann et al. 2014) are potential avenues. Second, despite the success of the regime forecasts, we did not consider the underlying forecast uncertainty in the economic application. A confidence measure for the probabilities could be useful for various applications, such as portfolio optimization or asset pricing. In this context, Alvarez et al. (2019) provide a foundation for future work. Finally, our approach can be extended to volatility forecasts and density forecasts. In addition, one could study international stock market indices or portfolios formed on industries or styles with the help of our forecasting model.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.ijforecast.2022.01.004.

References

Alvarez, R., Camacho, M., & Ruiz, M. (2019). Inference on filtered and smoothed probabilities in Markov-switching autoregressive models. *Journal Of Business And Economic Statistics*, *37*(3), 484–495.

Ang, A., & Bekaert, G. (2002). International asset allocation with regime shifts. *Review Of Financial Studies*, *15*(4), 1137–1187.

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

Bai, J., & Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, *70*(1), 191–221.

Bai, J., & Ng, S. (2008). Forecasting economic time series using targeted predictors. *Journal Of Econometrics*, *146*(2), 304–317.

Bekaert, G., & Hoerova, M. (2014). The VIX, the variance premium and stock market volatility. *Journal Of Econometrics*, *183*(2), 181–192.

Bollerslev, T., Tauchen, G., & Zhou, H. (2009). Expected stock returns and variance risk premia. *Review Of Financial Studies*, *22*(11), 4463–4492.

Brock, W., Lakonishok, J., & LeBaron, B. (1992). Simple technical trading rules and the stochastic properties of stock returns. *Journal Of Finance*, *47*(5), 1731–1764.

Çakmaklı, C., & van Dijk, D. (2016). Getting the most out of macroeconomic information for predicting excess stock returns. *International Journal Of Forecasting*, *32*(3), 650–668.

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

Campbell, J. Y., & Thompson, S. B. (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? *Review Of Financial Studies*, *21*(4), 1509–1531.

Campbell, J. Y., & Yogo, M. (2006). Efficient tests of stock return predictability. *Journal Of Financial Economics*, *81*(1), 27–60.

Chauvet, M., & Potter, S. (2000). Coincident and leading indicators of the stock market. *Journal Of Empirical Finance*, *7*(1), 87–111.

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

Chen, S.-S. (2012). Consumer confidence and stock returns over market fluctuations. *Quantitative Finance*, *12*(10), 1585–1597.

Chow, G., Jacquier, E., Kritzman, M., & Lowry, K. (1999). Optimal portfolios in good times and bad. *Financial Analysts Journal*, *55*(3), 65–73.

Clark, T. E., & West, K. D. (2007). Approximately normal tests for equal predictive accuracy in nested models. *Journal Of Econometrics*, *138*(1), 291–311.

Coudert, V., & Gex, M. (2008). Does risk aversion drive financial crises? Testing the predictive power of empirical indicators. *Journal Of Empirical Finance*, *15*(2), 167–184.

Cremers, K. M. (2002). Stock return predictability: A Bayesian model selection perspective. *Review Of Financial Studies*, *15*(4), 1223–1249.

Dal Pra, G., Guidolin, M., Pedio, M., & Vasile, F. (2018). Regime shifts in excess stock return predictability: An out-of-sample portfolio analysis. *Journal Of Portfolio Management*, *44*(3), 10–24.

DeLong, E. R., DeLong, D. M., & Clarke-Pearson, D. L. (1988). Comparing the areas under two or more correlated receiver operating characteristic curves: A nonparametric approach. *Biometrics*, *44*(3), 837–845.

Diebold, F. X., Lee, J.-H., & Weinbach, G. C. (1994). Regime switching with time-varying transition probabilities. *Business Cycles: Durations, Dynamics, And Forecasting*, *1*, 144–165.

Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal Of Business And Economic Statistics*, *13*(3), 253–263.

Erichson, N. B., Zheng, P., Manohar, K., Brunton, S. L., Kutz, J. N., & Aravkin, A. Y. (2020). Sparse principal component analysis via variable projection. *SIAM Journal On Applied Mathematics*, *80*(2), 977–1002.

Estrella, A., & Mishkin, F. S. (1996). The yield curve as a predictor of US recessions. *Current Issues In Economics And Finance*, *2*(7), 1–6.

Estrella, A., & Mishkin, F. S. (1998). Predicting US recessions: Financial variables as leading indicators. *Review Of Economics And Statistics*, *80*(1), 45–61.

Fama, E. F., & French, K. R. (1988). Dividend yields and expected stock returns. *Journal Of Financial Economics*, *22*(1), 3–25.

Fama, E. F., & French, K. R. (1989). Business conditions and expected returns on stocks and bonds. *Journal Of Financial Economics*, *25*(1), 23–49.

Fawcett, T. (2006). An introduction to ROC analysis. *Pattern Recognition Letters*, *27*(8), 861–874.

Focardi, S. M., Fabozzi, F. J., & Mazza, D. (2019). Modeling local trends with regime shifting models with time-varying probabilities. *International Review Of Financial Analysis*, *66*, Article 101368.

Gonzalez, L., Hoang, P., Powell, J. G., & Jing, S. (2006). Defining and dating bull and bear markets: Two centuries of evidence. *Multinational Finance Journal*, *10*(1/2), 81–116.

Guidolin, M., & Hyde, S. (2012). Can VAR models capture regime shifts in asset returns? A long-horizon strategic asset allocation perspective. *Journal Of Banking And Finance*, *36*(3), 695–716.

Guidolin, M., & Pedio, M. (2020). *Distilling large information sets to forecast commodity returns: Automatic variable selection or hidden Markov models?, Vol. 20140*: *BAFFI CAREFIN Working papers*.

Guidolin, M., & Pedio, M. (2021). Forecasting commodity futures returns with stepwise regressions: Do commodity-specific factors help? *Annals Of Operations Research*, *299*(1), 1317–1356.

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

Haas, M., Mittnik, S., & Paolella, M. S. (2004). A new approach to Markov-switching GARCH models. *Journal Of Financial Econometrics*, *2*(4), 493–530.
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.
Hamilton, J. D. (1994). *Time series analysis, Vol. 2*. Princeton University Press Princeton, NJ.
Hamilton, J. D., & Lin, G. (1996). Stock market volatility and the business cycle. *Journal Of Applied Econometrics*, *11*(5), 573–593.
Hammerschmid, R., & Lohre, H. (2018). Regime shifts and stock return predictability. *International Review Of Economics And Finance*, *56*(C), 138–160.
Hansen, B. E. (1991). The likelihood test under non-standard conditions: Testing the Markov trend model of GNP. *Journal Of Applied Econometrics*, *7*, 61–82.
Hauptmann, J., Hoppenkamps, A., Min, A., Ramsauer, F., & Zagst, R. (2014). Forecasting market turbulence using regime-switching models. *Financial Markets And Portfolio Management*, *28*(2), 139–164.
Henkel, S. J., Martin, J. S., & Nardari, F. (2011). Time-varying short-horizon predictability. *Journal Of Financial Economics*, *99*(3), 560–580.
Hong, H., Torous, W., & Valkanov, R. (2007). Do industries lead stock markets? *Journal Of Financial Economics*, *83*(2), 367–396.
Johnson, M. A., & Watson, K. J. (2011). Can changes in the purchasing managers' index foretell stock returns? an additional forward-looking sentiment indicator. *Journal Of Investing*, *20*(4), 89–98.
Kaya, H., Lee, W., & Pornrojnangkool, B. (2010). Regimes: Nonparametric identification and forecasting. *Journal Of Portfolio Management*, *36*(2), 94–105.
Kole, E., & van Dijk, D. (2017). How to identify and forecast bull and bear markets? *Journal Of Applied Econometrics*, *32*(1), 120–139.
Kristensen, J. T. (2017). Diffusion indexes with sparse loadings. *Journal Of Business And Economic Statistics*, *35*(3), 434–451.
Kritzman, M., & Li, Y. (2010). Skulls, financial turbulence, and risk management. *Financial Analysts Journal*, *66*(5), 30–41.
Kritzman, M., Li, Y., Page, S., & Rigobon, R. (2011). Principal components as a measure of systemic risk. *Journal Of Portfolio Management*, *37*(4), 112–126.
Kritzman, M., Page, S., & Turkington, D. (2012). Regime shifts: Implications for dynamic strategies (corrected). *Financial Analysts Journal*, *68*(3), 22–39.
Lewellen, J. (2004). Predicting returns with financial ratios. *Journal Of Financial Economics*, *74*(2), 209–235.
Ludvigson, S. C., & Ng, S. (2007). The empirical risk–return relation: A factor analysis approach. *Journal Of Financial Economics*, *83*(1), 171–222.
Ludvigson, S. C., & Ng, S. (2009). Macro factors in bond risk premia. *Review Of Financial Studies*, *22*(12), 5027–5067.
Lunde, A., & Timmermann, A. (2004). Duration dependence in stock prices: An analysis of bull and bear markets. *Journal Of Business And Economic Statistics*, *22*(3), 253–273.
Maheu, J. M., & McCurdy, T. H. (2000). Identifying bull and bear markets in stock returns. *Journal Of Business And Economic Statistics*, *18*(1), 100–112.
Maheu, J. M., McCurdy, T. H., & Song, Y. (2012). Components of bull and bear markets: Bull corrections and bear rallies. *Journal Of Business And Economic Statistics*, *30*(3), 391–403.
Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *Review Of Economics And Statistics*, *51*(3), 247–257.
Mönch, E. (2008). Forecasting the yield curve in a data-rich environment: A no-arbitrage factor-augmented VAR approach. *Journal Of Econometrics*, *146*(1), 26–43.
Neely, C. J., Rapach, D. E., Tu, J., & Zhou, G. (2014). Forecasting the equity risk premium: The role of technical indicators. *Management Science*, *60*(7), 1772–1791.
Pagan, A. R., & Sossounov, K. A. (2003). A simple framework for analysing bull and bear markets. *Journal Of Applied Econometrics*, *18*(1), 23–46.
Pástor, L., & Stambaugh, R. F. (2001). The equity premium and structural breaks. *Journal Of Finance*, *56*(4), 1207–1239.
Pesaran, M. H., & Timmermann, A. (1995). Predictability of stock returns: Robustness and economic significance. *Journal Of Finance*, *50*(4), 1201–1228.
Pettenuzzo, D., & Timmermann, A. (2011). Predictability of stock returns and asset allocation under structural breaks. *Journal Of Econometrics*, *164*(1), 60–78.
Rapach, D. E., Strauss, J. K., & Zhou, G. (2010). Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *Review Of Financial Studies*, *23*(2), 821–862.
Rapach, D. E., Wohar, M. E., & Rangvid, J. (2005). Macro variables and international stock return predictability. *International Journal Of Forecasting*, *21*(1), 137–166.
Rapach, D., & Zhou, G. (2013). Forecasting stock returns. In G. Elliot, & A. Timmermann (Eds.), *Handbook of economic forecasting, Vol. 2A* (pp. 328–383). Amsterdam: Elsevier.
Rapach, D., & Zhou, G. (2019). *Sparse macro factors*: *Working paper*, http://dx.doi.org/10.2139/Ssrn.3259447.
Rubbaniy, G., Asmerom, R., Rizvi, S. K. A., & Naqvi, B. (2014). Do fear indices help predict stock returns? *Quantitative Finance*, *14*(5), 831–847.
Sander, M. (2018). Market timing over the business cycle. *Journal Of Empirical Finance*, *46*(C), 130–145.
Schaller, H., & Norden, S. V. (1997). Regime switching in stock market returns. *Applied Financial Economics*, *7*(2), 177–191.
Timmermann, A. (2006). Forecast combinations. In G. Elliott, C. Granger, & A. Timmermann (Eds.), *Handbook of economic forecasting, Vol. 1* (pp. 135–196). Amsterdam: Elsevier.
Zens, G., & Böck, M. (2019). *A factor-augmented Markov switching (FAMS) model*: *Working paper*, https://Arxiv.Org/Abs/1904.13194.
Zhu, X., & Zhu, J. (2013). Predicting stock returns: A regime-switching combination approach and economic links. *Journal Of Banking And Finance*, *37*(11), 4120–4133.
Zou, H., Hastie, T., & Tibshirani, R. (2006). Sparse principal component analysis. *Journal Of Computational And Graphical Statistics*, *15*(2), 265–286.