UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

Dipartimento di Economia Marco Biagi

DEMB Working Paper Series

N. 259

Investor Sentiment Indicators and the Prediction of European Equity Index Returns:

A Machine Learning Approachand team in organizations

Carlos Alberto Barua Mimbela[1], Silvia Muzzioli[2]

May 2025

[1] University of Modena and Reggio Emilia, Department of Economics Marco Biagi
Email: carlos.barua@unimore.it
[1] University of Modena and Reggio Emilia, Department of Economics Marco Biagi and CEFIN
Email: silvia.muzzioli@unimore.it

ISSN: 2281-440X online

Investor Sentiment Indicators and the Prediction of European Equity Index Returns: A Machine Learning Approach

Carlos Alberto Barua Mimbela[1], Silvia Muzzioli[2]

Abstract

This paper investigates the predictive power of various sentiment indicators on European equity index returns, focusing on the EUROSTOXX600 and EUROSTOXX50 over the period 2010–2022. Using a set of sentiment indicators ranging from options-based measures (Call/Put ratios, Skew), economic and political uncertainty (EPU, RAX), and sector-based sentiment proxies (EASS indices), we estimate a suite of models from classical linear regression to regularized models (Ridge, Lasso) and machine learning techniques (Random Forest, Gradient Boosting). The aim is to assess both the statistical significance and predictive importance of these sentiment measures. Our findings show that sentiment indicators carry significant explanatory power for European equity returns, with certain indicators consistently emerging as relevant across methodologies. The results have implications for asset managers and policymakers interested in incorporating behavioral factors into asset pricing models.

Keywords*: Investor Sentiment, Machine Learning, European Equity Markets, Asset Pricing, Sentiment Indicators*

[1] Marco Biagi Department of Economics, University of Modena and Reggio Emilia, Italy.
Email: carlos.barua@unimore.it
[2] Marco Biagi Department of Economics and CEFIN, University of Modena and Reggio Emilia, Italy.
Email: silvia.muzzioli@unimore.it

1. Introduction

Investor sentiment can be defined as expectations about future cash flows and investment risks that are not fully justified by available fundamentals, or as a general mood of optimism or pessimism about financial markets. Sentiment plays a crucial role in the decision-making process of investors and can influence asset prices, particularly during periods of excessive exuberance or unjustified panic. Psychological factors, such as overconfidence, loss aversion, and herd behavior, often interact with sentiment, generating dynamics that deviate from the predictions of traditional rational asset pricing models. As a result, sentiment has emerged as a key behavioral component in modern finance theory, and a growing body of empirical literature has explored its relevance for predicting future returns, volatility, and co-movement across asset classes.

While much of the existing literature on investor sentiment is centered on the U.S. market and often leverages firm-level sentiment scores derived from textual analysis or media content, relatively few studies have focused on the European context or on aggregate indices such as the EUROSTOXX600 and EUROSTOXX50. Moreover, most studies rely on a narrow set of sentiment measures or assume linear relationships between sentiment and returns, potentially missing the complex interactions and nonlinearities that characterize financial markets. In particular, the influence of sectoral sentiment indicators and economic uncertainty measures has been underexplored in the context of pan-European indices.

This study aims to fill this gap by analyzing the impact of a broad and diverse set of sentiment indicators on the monthly returns of the EUROSTOXX600 and EUROSTOXX50 over the period 2010–2022. Specifically, we consider a comprehensive list of sentiment proxies, including options market indicators (Call/Put ratios, Skew), economic and political risk indices (EPU, RAX), sovereign credit spreads (SOVCISS), and a wide range of sector-specific sentiment measures (EASS indices). These indicators collectively capture different dimensions of market psychology, from investor hedging behavior and risk aversion to macroeconomic and geopolitical uncertainty.

To evaluate the predictive ability of these sentiment indicators, we adopt a time-series regression framework, combining traditional econometric models such as Ordinary Least Squares (OLS) with modern statistical learning techniques. In particular, we implement Ridge and Lasso regressions to address issues of multicollinearity and to perform regularized variable selection. Furthermore, we apply ensemble machine learning models, including Random Forests and Gradient Boosting Machines, to capture potential nonlinear effects and variable interactions that are not easily detected by linear models.

The use of machine learning in this context is motivated by its capacity to manage high-dimensional data and to uncover hidden patterns in the relationships between sentiment and market returns. Unlike conventional models, tree-based methods allow for flexible functional forms and are robust to overfitting when appropriate cross-validation techniques are employed. By comparing performance across models, we are able to assess the robustness and relevance of different sentiment indicators under both linear and nonlinear assumptions.

Our empirical findings indicate that several sentiment indicators—notably EPU, RAX, SOVCISS, and selected EASS indices—exhibit robust predictive power for European index returns. Moreover, we find that machine learning models outperform linear regressions in terms of out-of-sample forecasting accuracy, underscoring the usefulness of non-linear models in capturing the complex dynamics of investor sentiment. These results carry important implications for both institutional investors seeking to enhance their asset allocation strategies and policymakers interested in monitoring systemic risk through behavioral channels.

The remainder of the paper is structured as follows: Section 2 reviews the relevant literature on sentiment and asset pricing; Section 3 details the data sources and econometric methodology; Section 4 presents the empirical findings and model comparisons; and Section 5 concludes with a discussion of policy implications and directions for future research.

2. Literature review

Investor sentiment has been broadly recognized as a non-fundamental driver of financial market behavior. It is commonly defined as the beliefs about future cash flows and investment risks that are not justified by observable facts, or more generally as the overall mood of optimism or pessimism that influences trading behavior (Baker and Wurgler, 2006). A growing body of research has shown that sentiment can play a pivotal role in financial markets, particularly during periods of market stress, exuberance, or uncertainty (Reis and Pinho, 2020).

Despite its relevance, investor sentiment is not directly observable, and therefore researchers have adopted various indirect approaches to construct sentiment indicators. The literature typically distinguishes three main methodologies. The first approach relies on composite indices that aggregate multiple financial and economic variables believed to reflect market sentiment. These proxy-based indices often combine variables such as trading volume, market volatility, dividend premiums, and IPO activity. For instance, Sun et al. (2016) and Xu and Zhou (2018) outline strategies to extract latent sentiment factors from such data. In the European context, one notable example is the EURsent index developed by Reis and Pinho (2020). The second approach is survey-based and builds sentiment measures from the responses of firms, investors, or analysts to standardized questionnaires. Well-known examples include the Economic Sentiment Indicator (ESI) and the ZEW Indicator of Economic Sentiment. Although widely used, survey-based indices are typically published at low frequency (monthly or quarterly) and may lag behind rapid developments in financial markets (González-Sánchez and Morales de Vega, 2021). A third, and more recent, approach involves the extraction of sentiment from textual data, including news articles, microblogging platforms, and internet search trends. Examples include the Economic Policy Uncertainty (EPU) Index by Baker et al. (2016), the FEARS index by Da et al. (2015), sentiment indices based on Financial Times reports (Huang et al., 2019), and Google search-based indices (Gao et al., 2020; Anastasiou and Drakos, 2021). These indicators offer higher frequency and responsiveness but often suffer from limitations such as the inability to distinguish between positive and negative sentiment. Despite these advances, sentiment research in the European context remains underdeveloped. Many existing indices are backward-looking or do not adequately reflect asymmetries in investor perception of good and bad news. Option-based measures like the VSTOXX are commonly used but fail to capture directional sentiment. Moreover, macro-financial sentiment indicators tend to neglect key elements such as sovereign debt concerns, which have been highly relevant in recent European crises. There is also no consensus on which variables should be included in sentiment indices or on the optimal aggregation techniques to construct them. Empirical evidence on the predictive power of sentiment indicators in Europe is mixed. Studies such as Reis and Pinho (2020), Baker et al. (2012), and Corredor et al. (2013) have explored the relationship between sentiment and returns, often reaching divergent conclusions. Additionally, very few studies have addressed the role of sentiment in predicting returns at the level of broad indices, as opposed to individual stocks or portfolios sorted by firm characteristics.

In sum, the literature highlights several gaps: the backward-looking nature of most EU sentiment indicators; their limited capacity to account for asymmetric effects; insufficient integration of sovereign risk measures; the lack of distinction between positive and negative sentiment; disagreement on variable selection and aggregation methods; and a general scarcity of robust empirical evidence on their predictive performance. This paper aims to address these issues by incorporating a broad range of sentiment proxies and applying advanced econometric and machine learning techniques to evaluate their role in forecasting European equity index returns. Building upon the identified gaps in the existing literature, this study pursues five main objectives aimed at enhancing both the conceptual and empirical understanding of investor sentiment in the European equity market.

First, the research seeks to develop forward-looking sentiment indices (FLSIs) constructed from option-based data. These indices are designed to incorporate anticipatory market signals and to differentiate between expectations of positive and negative returns, thus addressing the commonly observed backward-looking nature of conventional sentiment measures and their inability to account for asymmetries in investor perceptions. Second, the study aims to design backward-looking sentiment indices (BLSIs) based on macro-financial variables and survey-based indicators, with an explicit focus on embedding information related to sovereign debt sustainability. This aspect is particularly relevant given the role of debt-related concerns in shaping sentiment during European financial crises and directly tackles the literature's neglect of such dimensions. Third, the project includes the construction of text-based sentiment indices (TSIs), utilizing data derived from social media platforms, microblogging services, and internet search behavior. The goal is to generate high-frequency sentiment indicators that are capable of distinguishing between good and bad news, thereby improving the interpretability and informativeness of sentiment proxies derived from unstructured data sources. Fourth, the study intends to integrate the various sentiment measures—FLSIs, BLSIs, and TSIs—into a comprehensive Market Sentiment Index (MSI). This process involves the careful selection of relevant variables and the adoption of aggregation techniques that respect the ordinal and heterogeneous nature of the underlying data, while also managing the uncertainty inherent in sentiment-based knowledge representation. This unified index will provide a multidimensional and robust measure of market sentiment suitable for empirical application.

Finally, the project aims to empirically test the information content and predictive power of the individual and combined sentiment indicators—FLSI, BLSI, TSI, and MSI—with respect to future returns of European equity indices. This analysis will allow us to assess the usefulness of such indicators for investors in terms of return forecasting and for policymakers in anticipating periods of financial stress, thus contributing new empirical evidence to a research area where existing findings remain limited and often inconclusive.

Table 1: *Literature review*

Paper	**Objective**	**Data Used**
Baker & Wurgler (2006)	To define and measure investor sentiment as a non-fundamental factor affecting financial markets.	Constructed a sentiment index using a set of financial market proxies, including trading volume, market volatility, dividend premiums, and IPO activity.
Reis & Pinho (2020)	To develop the EURsent index capturing investor sentiment in the European context.	Macro-financial indicators specific to European markets, aggregated into a composite sentiment measure.
Sun et al. (2016)	To extract latent sentiment factors from market data using factor analysis techniques.	Financial market variables such as volatility, trading volume, and dividend-related indicators.

Xu & Zhou (2018)	To propose methods for constructing sentiment indices from market-based data.	Composite indicators based on observable financial variables reflecting investor behavior.
González-Sánchez & Morales de Vega (2021)	To critically assess the limitations of survey-based sentiment indicators.	Survey data from sources such as the Economic Sentiment Indicator (ESI) and the ZEW Indicator; typically low-frequency and potentially lagging indicators.
Baker et al. (2016)	To construct the Economic Policy Uncertainty (EPU) Index as a proxy for uncertainty in financial markets.	Textual analysis of newspaper articles using keyword frequency related to policy uncertainty.
Da et al. (2015)	To develop the FEARS Index as a measure of investor anxiety based on online behavior.	Google Search Volume Index (GSVI) data capturing the frequency of specific fear-related terms.
Huang et al. (2019)	To extract sentiment indices from financial journalism content.	Textual sentiment analysis of *Financial Times* articles using natural language processing techniques.
Gao et al. (2020)	To build sentiment proxies using search engine data as a reflection of investor attention.	Google Trends data related to financial and economic topics.
Anastasiou & Drakos (2021)	To assess the predictive power of internet-based sentiment indicators in financial markets.	Internet search trends and behavioral indicators based on user activity on digital platforms.
Baker et al. (2012)	To empirically test the relationship between sentiment and future stock returns.	Sentiment indices and U.S. stock market return data.
Corredor et al. (2013)	To examine the impact of sentiment on asset returns within the European market.	European sentiment indicators and financial return data across equity markets.

3. Data and Methodology

This study investigates the predictive role of investor sentiment on European equity markets by focusing on the monthly returns of two key stock indices: the EUROSTOXX600 and the EUROSTOXX50. The EUROSTOXX600 provides a broad representation of the European equity universe, encompassing 600 companies from 17 European countries across large, mid, and small capitalizations. It is widely used as a benchmark for diversified European portfolios and reflects the overall economic and financial conditions of the continent. The EUROSTOXX50, by contrast, comprises the 50 largest and most liquid companies within the Eurozone and is commonly interpreted as a measure of core Euro area performance, particularly in relation to large-cap blue-chip firms. Analyzing both indices allows for an assessment of sentiment effects at different levels of market aggregation and capital intensity.

The dataset used for the empirical analysis was obtained from the Refinitiv database and covers the period from January 2010 to December 2022. Monthly closing prices for both indices were transformed into log returns, computed as the natural logarithm of the ratio between the index level at time $t+1$ and at time t. This transformation yields continuously compounded returns, which offer desirable statistical properties such as time additivity and approximate normality. These features are particularly useful in linear and nonlinear modeling frameworks where variance stabilization and symmetry of the distribution improve estimation and inference. To investigate whether sentiment contains information relevant for explaining and predicting equity index performance, a comprehensive set of sentiment indicators was collected and prepared for analysis. The explanatory variables include measures derived from options markets, such as call and put volumes as well as the skew index, which reflects perceived asymmetry in tail risk. Other variables capture economic and policy-related uncertainty, such as the RAX index and the Economic Policy Uncertainty (EPU) index. Sovereign risk is proxied through the SOVCISS index, which measures stress in Eurozone sovereign debt markets. Additionally, sentiment measures aggregated at the sectoral level—such as those from the European Aggregate Sentiment Scores (EASS)—offer insight into investor perceptions across different market segments, including commodities, corporate equities, bonds, money markets, financial institutions, foreign exchange, and cross-asset interactions. All sentiment variables were standardized to ensure comparability and to mitigate issues related to scale heterogeneity. The methodological approach follows a two-step structure. Initially, a classical linear regression model is estimated using ordinary least squares (OLS), with the log returns of the EUROSTOXX600 and EUROSTOXX50 serving as dependent variables and the sentiment indicators as explanatory variables. This provides a baseline assessment of the linear relationships and the statistical significance of each sentiment measure. However, linear models may fail to capture the complexity of financial relationships, especially in the presence of multicollinearity or nonlinear interactions among predictors. To address these limitations, the analysis proceeds by employing a range of more flexible modeling techniques drawn from the machine learning literature. Among these, ridge regression is adopted to address multicollinearity by introducing a penalty term that shrinks coefficient estimates while retaining all variables in the model. Lasso regression, in contrast, performs variable selection by applying a penalty that can reduce some coefficients exactly to zero, thereby enhancing model interpretability and focusing attention on the most informative predictors. To capture potential nonlinearities and interactions, ensemble methods such as random forests and gradient boosting are utilized. Random forests build multiple decision trees using bootstrap samples and average their predictions to reduce variance and overfitting. Gradient boosting constructs a sequence of trees where each successive model improves upon the errors of the previous one, allowing for highly accurate and fine-tuned predictions even in complex data environments.

All models are estimated under rigorous cross-validation procedures to ensure the robustness and generalizability of the results. In-sample model fit is assessed using metrics such as the coefficient of determination (R^2) and its adjusted version, while out-of-sample performance is evaluated based on root mean squared error (RMSE) and mean absolute error (MAE). Furthermore, for the nonlinear models, the relative importance of each predictor is quantified through variable importance scores, offering insights into the most influential sentiment indicators in forecasting European equity index returns. This multi-method approach is intended to provide a comprehensive evaluation of the role of sentiment in European markets, combining interpretability and flexibility to capture both traditional linear effects and more complex nonlinear dynamics.

4. Analysis of results

Linear on EUROSTOXX600 returns

The empirical analysis begins with an exploratory investigation of the relationship between sentiment indicators and monthly returns of the EUROSTOXX600 index. A correlation matrix was computed to assess the linear dependence between the regressors and the dependent variable. Among the sentiment variables considered, the strongest correlations with EUROSTOXX600 returns were observed for the call and put volumes, as well as the EASSCROS index, with correlation coefficients around ±0.20. While modest in absolute value, these correlations suggest a potentially meaningful—albeit partial—linear association, which warranted further investigation through regression models.

To this end, a series of linear regressions were estimated using different subsets of sentiment indicators. The first model includes the three variables most correlated with returns: call, put, and EASSCROS. The estimation results indicate that both call and put are statistically significant at conventional levels. In particular, the call variable is negatively associated with returns (coefficient ≈ -0.013, $p < 0.001$), while the put variable exhibits a positive and weaker relationship (coefficient ≈ 0.0011, $p \approx 0.015$). These findings may be interpreted as follows: a higher call volume, possibly associated with optimistic positioning, might reflect overconfidence that precedes market corrections, whereas put volume, commonly interpreted as a hedge against downside risk, may signal increased caution that precedes positive performance. The EASSCROS variable, in contrast, is not statistically significant in this specification. The model explains approximately 13% of the total variance in index returns ($R^2 = 0.131$), which is a non-negligible figure considering the inherent noise in financial return data. In a second regression, the model is extended to include several additional EASS indicators (commodities, bonds, financials, equities, foreign exchange, and cross-asset signals). While the call and put variables remain significant, the EASS variables do not display statistical relevance in this broader setting. The coefficients associated with sectoral sentiment proxies are generally small and accompanied by high *p*-values, suggesting that their individual explanatory power is limited or possibly masked by multicollinearity or omitted interactions. The model's R^2 increases slightly to 0.152, but the adjusted R^2—which accounts for the number of regressors—declines to 0.106, indicating that the inclusion of these additional variables may lead to overfitting without substantial gain in explanatory power. A third and more comprehensive regression includes a broader set of predictors, incorporating macro-financial uncertainty indices (RAX and EPU), skewness, sovereign credit risk (SOVCISS), and all EASS sentiment scores. In this specification, several noteworthy patterns emerge. First, the EASSCOMP, EASSBOND, EASSMKT, EASSFINI, EASSEQUI, EASSFORE, and EASSCROS variables all become highly statistically significant, with very similar positive coefficient, suggesting a strong and homogeneous influence of sectoral sentiment on market returns when jointly considered. The high degree of similarity in their estimates—each with coefficients close to 2.12—may reflect shared information content across EASS dimensions, or possibly a common latent sentiment factor embedded in their construction. Second, the SOVCISS index is also significant, with a negative coefficient ($\approx$ -0.00145, $p < 0.01$), indicating that increased sovereign credit stress is associated with lower future index returns. Interestingly, the RAX and EPU indices, as well as skew, do not appear to significantly contribute in this linear framework. The model's explanatory power improves notably in this specification, with an R^2 of approximately 0.246 and an adjusted R^2 of 0.171.

Overall, these results highlight several insights. Simpler models including only the most correlated sentiment variables already reveal statistically significant relationships with market returns. However, when a broader sentiment framework is introduced, including both macro risk perceptions and sectoral dynamics, the explanatory power increases substantially. Notably, EASS indicators show strong significance only when jointly included, suggesting interaction or complementary effects among sentiment channels. At the same time, multicollinearity may obscure the individual effects of certain predictors, motivating the use of regularization and nonlinear modelling, as discussed in the subsequent section.

Linear on EUROSTOXX50 returns

The analysis is extended to the EUROSTOXX50, an index composed of the 50 largest and most liquid companies in the Euro area. As with the broader index, the correlation analysis reveals that call, put, and EASSCROS are the most linearly correlated with returns, again with coefficients close to ±0.20.

The first regression model, focusing on these three variables, shows that both call and put remain statistically significant predictors, with a negative sign for call and a positive sign for put. The EASSCROS variable is again not significant in this reduced specification. The model explains 11.4% of the variance in EUROSTOXX50 returns ($R^2 = 0.114$), slightly lower than the corresponding value for the broader index, suggesting a slightly weaker explanatory power when applied to blue-chip stocks. In the second regression, where additional EASS variables are included, the call and put effects persist while none of the sectoral sentiment proxies appear to add explanatory power individually. The model's R^2 increases to 0.133, but the adjusted R^2 decreases to 0.086, indicating that the marginal gain in fit does not justify the added model complexity. As with the EUROSTOXX600, the third and most comprehensive specification produces the richest results. In this case, a consistent pattern of statistically significant coefficients emerges across all EASS indicators. The variables EASSCOMP, EASSBOND, EASSMMKT, EASSFINI, EASSEQUI, EASSFORE, and EASSCROS all exhibit positive and significant effects. As in the previous model, the homogeneity in coefficient magnitude suggests a strong shared component. Additionally, the SOVCISS index is again negatively associated with returns and statistically significant, indicating that sovereign credit stress is a relevant factor for both broad and concentrated equity benchmarks. Other macro indicators such as RAX, EPU, and skew remain statistically insignificant. This model achieves an R^2 of 0.206 and an adjusted R^2 of 0.127. Comparing results across the two indices reveals a consistent pattern: the call and put variables hold predictive value in simple models, while the inclusion of a broader sentiment structure only improves explanatory power when the entire network of EASS indicators is modeled simultaneously. Notably, the explanatory power is slightly higher for the EUROSTOXX600 than for the EUROSTOXX50, possibly due to the more diverse composition of the former index, which may better reflect heterogeneous investor sentiment across sectors and firm sizes.

These findings provide compelling evidence that sentiment-based variables, especially when considered in an integrated manner, possess predictive power for European stock market returns. They also suggest that sector-specific sentiment channels are more informative when analyzed jointly rather than in isolation, and that sovereign risk remains a key variable in European equity pricing.

Non-Linear Models: Lasso Regression (EUROSTOXX600)

To complement the results obtained from classical linear models, we apply Lasso regression to assess the predictive power of sentiment indicators for the EUROSTOXX600 index. Lasso (Least Absolute Shrinkage and Selection Operator) is a regularization technique that adds an L1 penalty to the regression objective function, encouraging sparsity in the estimated coefficients. This approach not only helps to address multicollinearity, but also performs automatic variable selection by shrinking the coefficients of less informative predictors to exactly zero. In doing so, Lasso provides a more parsimonious model and helps identify the most relevant predictors, particularly in settings with a high number of correlated variables.

The model is estimated using cross-validation to determine the optimal penalty parameter (lambda), which in this case is approximately 0.00064. The Lasso regression retains a subset of predictors with non-zero coefficients, while setting others to zero, indicating their limited explanatory power in the presence of more informative variables. The results confirm the importance of the call and put variables, both of which retain significant weights, with call having a negative coefficient (≈ -0.0103) and put a positive one (≈ 0.0018). This is consistent with the findings from the OLS models and reinforces the robustness of these two variables across methodologies. Additionally, RAX is retained in the model with a moderate negative coefficient, suggesting that heightened risk aversion is associated with lower future returns. Among the sectoral sentiment indicators, EASSMMKT, EASSFINI, and EASSCROS are preserved, indicating some predictive relevance for money markets, financial institutions, and cross-asset sentiment, respectively. The coefficients are generally small in magnitude, but their persistence across model specifications suggests a degree of robustness. Interestingly, EASSEQUI and EASSFORE enter the model with negative weights, possibly capturing pessimism in equity and currency markets that translates into weaker future returns. The SOVCISS index also remains in the model with a small negative coefficient, consistent with previous findings on the detrimental impact of sovereign risk on equity performance. In contrast, several variables—such as skew, EPU, and multiple EASS indicators including commodities, bonds, and general company sentiment—are eliminated by the Lasso penalty. Their exclusion suggests either low marginal predictive value or collinearity with more informative sentiment proxies.

Overall, the Lasso model achieves an R^2 of 0.188, slightly lower than the full OLS specification but with a higher interpretability due to reduced dimensionality. The adjusted R^2, which penalizes model complexity, is 0.108, suggesting that despite fewer active predictors, the model retains a substantial proportion of explanatory power. These results illustrate the usefulness of Lasso regression in identifying a compact subset of sentiment variables that contribute meaningfully to forecasting equity index returns in the European context.

Non-Linear Models: Lasso Regression (EUROSTOXX50)

The Lasso regression was also applied to the EUROSTOXX50 index in order to assess whether the set of sentiment indicators retains predictive power in the context of a more concentrated large-cap benchmark. As with the EUROSTOXX600, the model was estimated using cross-validation to determine the optimal regularization parameter, which in this case is approximately 0.00272. The relatively higher lambda value compared to the EUROSTOXX600 model indicates a stronger penalization, resulting in a more parsimonious solution.

The Lasso procedure selects a reduced subset of predictors, retaining only call, put, EASSEQUI, and EASSCROS, while setting the coefficients of all remaining variables—including skew, RAX, EPU, and most EASS sectoral indicators—to zero. This selection implies that for the EUROSTOXX50, fewer sentiment indicators contain incremental predictive information once penalization is applied. Among the retained variables, the call volume continues to display a negative coefficient ($\approx$ -0.0086), consistent with the hypothesis that increased bullish speculative activity may precede lower future returns, possibly due to over-optimism. The put variable is positively signed, albeit with a smaller magnitude ($\approx$ 0.00066), again confirming its potential role as a contrarian signal. The EASSEQUI index is retained with a small negative coefficient, possibly capturing sentiment deterioration within equity markets that translates into weaker future performance. Lastly, EASSCROS is included with a marginal positive weight (≈ 0.00021), suggesting that cross-asset sentiment, while weakly predictive, may still contribute to explaining variation in index-level returns. Importantly, several indicators that were significant in the OLS specifications—including EASSCOMP, EASSBOND, SOVCISS, and macroeconomic sentiment proxies such as EPU and RAX—are excluded by the Lasso model. Their removal indicates limited marginal contribution in the presence of the selected variables, or redundancy due to collinearity. The sparsity of the final model confirms that, for a concentrated index like the EUROSTOXX50, sentiment signals may be less dispersed across sectors and more difficult to detect outside core equity and options-based channels. From a performance standpoint, the Lasso model explains approximately 10.3% of the total variance ($R^2 = 0.103$), with a much lower adjusted R^2 of 0.014, suggesting that the gain in parsimony comes at the cost of predictive power. Compared to the Lasso model for the EUROSTOXX600, the explanatory capacity is notably weaker, reinforcing the hypothesis that broad-based indices may be more sensitive to a diversified set of sentiment signals, while large-cap indices respond to a narrower subset of investor expectations.

In sum, while Lasso regression confirms the relevance of certain sentiment measures—particularly call/put activity and equity-sector sentiment—it also highlights the more selective and limited nature of predictability in the EUROSTOXX50, emphasizing the importance of index composition in determining the efficacy of sentiment-based forecasting models.

Non-Linear Models: Ridge Regression (EUROSTOXX600)

To complement the variable selection properties of Lasso, the Ridge regression model is employed to assess the predictive contribution of the sentiment variables when multicollinearity is present but exclusion of variables is not desired. Ridge regression, unlike Lasso, applies an L2 penalty to the magnitude of coefficients, shrinking them toward zero without setting any of them exactly to zero. This approach is particularly effective in the presence of correlated predictors, as it stabilizes estimation and reduces overfitting while preserving all available explanatory information.

In this analysis, the sentiment indicators retained from the Lasso regression for the EUROSTOXX600 were used as the basis for the Ridge estimation. The optimal penalty parameter (lambda) was determined via cross-validation and found to be approximately 0.00139. The Ridge model yields a more stable set of coefficients compared to OLS, with all predictors included and appropriately regularized. The estimated coefficients reveal a consistent pattern with previous results. The call variable continues to exhibit a negative association with returns (≈ -0.00995), while the put variable maintains a small positive effect (≈ 0.00200). The RAX index enters with a negative sign, suggesting that rising market risk aversion is associated with declining future returns. The EPU index, although small in magnitude, is retained with a positive coefficient, potentially reflecting some lagged adjustment of macroeconomic uncertainty into equity prices. Several sectoral sentiment indicators display meaningful contributions: EASSMMKT, EASSFINI, and EASSCROS are positively signed, while EASSEQUI and EASSFORE are negatively associated with returns. These results mirror the qualitative patterns observed in the Lasso regression, but with more stable effect sizes across variables. The SOVCISS index again enters with a negative sign, reinforcing the idea that sovereign credit stress acts as a drag on market performance. From a performance standpoint, the Ridge model achieves an R^2 of 0.2005, with an adjusted R^2 of 0.1453, slightly higher than those obtained with Lasso. While Ridge does not perform variable selection, its ability to include all predictors without overfitting allows for a richer representation of sentiment dynamics. In particular, it confirms the relevance of option-based signals, risk aversion, and selected sector-specific sentiment indices in explaining variation in monthly returns of the EUROSTOXX600. These results suggest that Ridge regression is particularly useful when the objective is to preserve model completeness without sacrificing stability, offering a compromise between the parsimony of Lasso and the inclusiveness of OLS.

Non-Linear Models: Ridge Regression (EUROSTOXX50)

Following the Lasso estimation, which identified a parsimonious set of relevant predictors for the EUROSTOXX50, we apply Ridge regression to the same group of variables in order to obtain a more stable and inclusive estimation framework. As previously noted, Ridge regression introduces an L2 penalty that reduces coefficient variance in the presence of multicollinearity, without excluding variables entirely. This makes Ridge particularly suitable for financial datasets, where predictors often contain overlapping information and interactions may be difficult to specify ex ante. Using cross-validation, the optimal penalty parameter was found to be 0.01236, higher than the corresponding value for the EUROSTOXX600 model, reflecting a need for stronger shrinkage in this more concentrated index context. Ridge estimation retains all sentiment variables, assigning smaller and regularized coefficients.

The results show that call and put volumes retain their expected signs. The call coefficient is negative (≈ -0.00747), suggesting that excess call activity may be associated with future corrections, while the put coefficient remains positive (≈ 0.00102), consistent with a contrarian reading of downside protection demand. However, both coefficients are smaller in magnitude than those estimated for the broader index, potentially reflecting the more muted responsiveness of large-cap stocks to short-term sentiment shifts. The RAX and EPU variables remain in the model but exhibit minimal influence. RAX is weakly negative (≈ -0.0066), while EPU is close to zero, implying limited predictive contribution from these aggregate uncertainty measures when controlling for other sentiment variables. Sectoral sentiment indicators provide mixed results. While EASSMMKT and EASSFINI are positively signed, their coefficients are very small. In contrast, EASSEQUI and EASSFORE carry negative coefficients, suggesting that pessimism in the equity and currency domains may weigh on large-cap returns. The EASSCROS index maintains a small positive value, indicating a slight forward-looking signal from cross-asset sentiment. Lastly, SOVCISS, the sovereign risk index, remains negative but its effect is weaker than in previous models (≈ -0.00011), possibly due to the relatively higher resilience of EUROSTOXX50 constituents to sovereign debt concerns. Overall, the Ridge regression for the EUROSTOXX50 explains approximately 13.6% of the variation in monthly returns ($R^2 = 0.1358$), with an adjusted R^2 of 0.0762, indicating moderate explanatory power. Compared to the Lasso model, Ridge preserves a richer set of predictors and confirms the relevance of several sentiment dimensions, albeit with weaker marginal contributions.

In sum, Ridge estimation highlights that even when sentiment variables are retained collectively, their explanatory strength for the EUROSTOXX50 is more diffuse and modest than for the EUROSTOXX600. This reinforces the notion that large-cap indices are less sensitive to sectoral and cross-asset sentiment signals, and that their returns are influenced by a narrower channel of investor expectations—chiefly option activity and equity-specific sentiment.

Non-Linear Models: Random Forest (EUROSTOXX600)

To further explore the predictive relationship between sentiment indicators and equity market returns, we estimate a Random Forest model for the EUROSTOXX600 index. Random Forests are ensemble learning methods based on decision trees, where multiple trees are trained on bootstrapped subsamples of the data and their predictions are aggregated—typically by averaging in the case of regression. This approach is particularly effective in capturing complex nonlinear relationships and interactions among predictors without requiring prior specification of functional forms. Moreover, Random Forests provide measures of variable importance, which are useful for identifying the predictors that contribute most to reducing forecast error.

In the current application, the Random Forest model was trained using the same sentiment indicators previously analyzed in linear and penalized regressions. However, the model fails to achieve positive predictive power, with an out-of-sample R^2 of -0.114, indicating that the model performs worse than a naïve mean predictor. This result suggests that Random Forest, in this context, may be overfitting to the training data or failing to capture meaningful nonlinear structure in the relationship between sentiment variables and index returns. Possible explanations include the relatively small sample size, the low signal-to-noise ratio in monthly returns, and the high dimensionality relative to data length.

Despite the poor global predictive performance, the model's internal variable importance metrics offer insights into the relative contribution of each sentiment indicator. The most influential variable is EASSCROS, which shows the highest increase in mean squared error when permuted, followed by call, eassmmkt, and eassfini. These findings partially align with those from linear and penalized models, confirming the relevance of call volume and cross-asset sentiment. Interestingly, put and sovciss also emerge as moderately important, although their explanatory power appears weaker than in the Ridge and Lasso settings.

Conversely, variables such as epu, eassequi, and eassfore display very low importance scores, suggesting limited contribution to the reduction of prediction error in a nonlinear framework. This result is consistent with their limited statistical significance in previous models.

Overall, the Random Forest model fails to improve prediction accuracy for EUROSTOXX600 returns but provides useful corroboration of the central role played by option-based sentiment (particularly call) and selected sectoral or cross-asset sentiment indicators. These findings highlight the limitations of tree-based ensemble methods in low-frequency, low-signal financial time series data, and emphasize the importance of carefully balancing model complexity with sample size and structure.

Non-Linear Models: Random Forest (EUROSTOXX50)

The Random Forest methodology was also applied to the EUROSTOXX50 index to examine whether nonlinearities or interaction effects among sentiment indicators could yield predictive improvements for this large-cap equity benchmark. As in the case of the EUROSTOXX600, the model was trained using a full set of sentiment-based predictors, and its performance was evaluated using out-of-sample R^2.

The results indicate that the Random Forest model performs poorly, with an out-of-sample R^2 of -0.1358. This negative value suggests that the model not only fails to capture the structure of the data but actually performs worse than a constant mean model. These findings are consistent with those observed for the EUROSTOXX600 and underscore the challenges of using nonparametric, data-hungry methods such as Random Forests in the context of relatively short monthly financial time series. High-dimensionality, combined with a limited sample size and inherently noisy return dynamics, appears to limit the efficacy of this technique in a low-frequency equity return forecasting context. Nevertheless, the model's internal measure of variable importance provides meaningful insights into the structure of the predictors. The call and put variables emerge as the most influential sentiment indicators, with the highest increases in mean squared error when permuted. This aligns well with the results obtained in linear, Lasso, and Ridge models, reinforcing the robustness of options market sentiment as a leading predictor of index-level returns. The EASSCROS index also appears prominently among the most important variables, followed by RAX, EASSFINI, and EASSFORE, which contribute moderately to node purity and error reduction. The relevance of cross-asset sentiment and financial sector sentiment reflects findings from prior models, although their marginal impact remains small in absolute terms. Other indicators, such as EPU, SOVCISS, EASSEQUI, and EASSMMKT, show very limited variable importance. This may indicate that macro-level uncertainty and sector-specific sentiment contribute relatively less to return prediction in large-cap indices like the EUROSTOXX50, which may already internalize broad risk factors through the behavior of dominant, multinational firms.

In sum, while the Random Forest model fails to deliver predictive gains for the EUROSTOXX50, it does confirm the central role of options-based sentiment—particularly call and put volumes—as well as certain aspects of sectoral and cross-asset investor sentiment. However, the overall ineffectiveness of the model suggests that high-complexity machine learning methods may not be well suited to monthly financial forecasting tasks when sample sizes are limited and signal strength is weak.

Non-Linear Models: Gradient Boosting (EUROSTOXX600)

As a final machine learning approach, we apply Gradient Boosting Regression to model the relationship between sentiment indicators and monthly returns of the EUROSTOXX600. Gradient Boosting is an ensemble technique that builds predictive models in a sequential fashion, where each new tree is trained to correct the residuals of the previous model. This method is particularly powerful in capturing complex nonlinear relationships and subtle interaction effects that may be undetectable through linear or even bagged tree models like Random Forests. Its adaptive nature, combined with the ability to minimize overfitting through regularization, makes Gradient Boosting especially appealing in financial modeling when sufficient tuning and cross-validation are employed. The model performs extremely well in-sample, achieving an R^2 of 0.9996, which indicates an almost perfect fit to the training data. While such a result should be interpreted with caution—since it may suggest overfitting in the absence of strong out-of-sample validation—it nevertheless shows that the sentiment indicators collectively contain sufficient structure to replicate the observed returns with high fidelity in-sample. Beyond predictive accuracy, Gradient Boosting allows for a detailed analysis of feature importance based on three metrics: Gain, which measures the relative contribution of each feature to reducing error; Cover, which reflects the proportion of observations affected by a feature; and Frequency, which counts how often a variable is used in splits across all trees. The call variable emerges as the most important feature across all dimensions, contributing more than 13% of the total model gain. This finding is consistent with prior results from linear, Lasso, Ridge, and Random Forest models, and confirms the persistent influence of option-based optimism in forecasting equity returns. Put also ranks highly, with strong coverage and frequency scores, reinforcing its predictive role. Among the sectoral sentiment indicators, EASSCOMN, EASSFORE, and EASSBOND appear particularly influential, with high gain and moderate frequency. Their importance suggests that commodity market sentiment, currency expectations, and bond market mood all carry relevant information for forecasting equity returns—perhaps reflecting macro-financial interconnections that are better captured through nonlinear learning algorithms. Additionally, EPU and RAX, which had limited explanatory power in linear settings, show substantial gain values here, indicating that policy and risk-related uncertainty may influence returns through complex, possibly threshold-based mechanisms. Other variables such as EASSCROS, EASSFINI, EASSEQUI, and EASSCOMP also contribute meaningfully, albeit with smaller shares of total gain. The SOVCISS index and skew appear less central but still relevant, suggesting a broader diffusion of sentiment impact across multiple domains.

Overall, the Gradient Boosting results underscore the latent structure embedded in sentiment data and highlight the ability of flexible models to extract meaningful signal from otherwise noisy predictors. However, given the near-perfect in-sample fit, further validation is necessary to assess whether these insights generalize to out-of-sample data. Without such external confirmation, caution must be exercised in interpreting predictive superiority as genuine economic signal rather than overfitting.

Non-Linear Models: Gradient Boosting (EUROSTOXX50)

The Gradient Boosting regression model was also applied to the EUROSTOXX50 index, with the aim of uncovering potentially nonlinear relationships between sentiment indicators and monthly equity returns for large-cap firms within the Euro area. As discussed previously, Gradient Boosting constructs an ensemble of weak learners—typically decision trees—where each subsequent model corrects the residuals of the preceding one. This iterative refinement process makes Gradient Boosting particularly adept at capturing complex interaction effects and higher-order nonlinearities, which may be overlooked by traditional regression models.

The model yields an exceptionally high in-sample R^2 of 0.9997, even slightly exceeding the already strong result observed for the EUROSTOXX600. Such a near-perfect fit, while illustrative of the model's flexibility and expressiveness, raises concerns regarding potential overfitting and highlights the need for careful out-of-sample validation. Nonetheless, the internal structure of the model, as revealed by feature importance metrics, offers valuable insights. The call variable once again emerges as the most influential predictor, contributing approximately 13.6% to total model gain. This reconfirms the pivotal role of option-based optimism in explaining future market movements. The put volume follows closely, along with EASSEQUI, the sentiment index specific to equities, which is particularly relevant in the context of an index composed exclusively of large-cap stocks. The high importance of these three features across both Gain and Cover dimensions reflects consistent and substantive contributions to the model's accuracy. Interestingly, macro-level indicators such as RAX and EPU appear more prominently in the EUROSTOXX50 model compared to the EUROSTOXX600. RAX contributes approximately 9.3% to total gain, and EPU roughly the same, suggesting that large-cap equity performance may be more sensitive to investor perceptions of systemic risk and policy uncertainty—factors that often affect multinational corporations more directly. Among the sectoral sentiment measures, EASSCOMN, EASSFORE, and EASSBOND stand out for their relatively high gain and coverage, indicating that macro-financial channels such as commodity prices, exchange rate expectations, and bond market sentiment may influence the return dynamics of the EUROSTOXX50. Additional signals are provided by EASSFINI, EASSCROS, and EASSMMKT, though with more modest contributions. The SOVCISS index and skew, while included, show comparatively lower gain, reinforcing their marginal predictive role in the large-cap setting. These results are largely consistent with the findings from the EUROSTOXX600 model, but some differences emerge. Notably, sentiment signals related to macro uncertainty (RAX, EPU) and the equity sector (EASSEQUI) gain more prominence in the EUROSTOXX50 case, possibly due to the nature of the index, which aggregates systemically relevant and globally exposed firms. The structure of importance also appears more evenly distributed, suggesting a richer interaction between global sentiment factors and index-level performance.

In summary, the Gradient Boosting model for the EUROSTOXX50 delivers a highly accurate in-sample fit and reinforces the dominant role of call, put, and equity-specific sentiment. However, the inclusion of macro sentiment variables and the broader distribution of feature contributions point to a more diversified set of drivers for large cap returns. These results highlight the model's capacity to integrate heterogeneous sentiment signals, though the high explanatory power observed in-sample should be interpreted cautiously in the absence of further validation.

Table 2. *Comparison of results*

Model	**R^2(STOXX600)**	**R^2(STOXX50)**	**Principal variables**
OLS (2 var)	0.131	0.114	call (–), put (+)
OLS (complete)	0.246	0.206	EASSCOMP (+), EASSBOND (+), EASSMMKT (+), SOVCISS (–), call
Lasso	0.188	0.103	call (–), put (+), RAX (–), EASSMMKT (+), EASSFINI (+), EASSCROS (+), SOVCISS (–) / only call, put, EASSCROS, EASSEQUI in the STOXX50
Ridge	0.201	0.136	call, put, RAX, EPU, EASSMMKT, EASSFINI, EASSFORE, EASSCROS, SOVCISS
Random Forest	–0.114	–0.136	call, put, EASSCROS, RAX, SOVCISS
Gradient Boosting	0.9996	0.9997	call, put, EPU, RAX, EASSCROS, EASSEQUI, EASSCOMN, EASSFORE, EASSBOND

5. Conclusion

This paper explores the relationship between investor sentiment and the monthly returns of two major European equity indices—EUROSTOXX600 and EUROSTOXX50—over the period 2010–2022. By employing a wide array of sentiment indicators, ranging from option-based measures and policy uncertainty indices to sector-specific sentiment scores, we aim to identify the most informative predictors of future market performance and assess their stability across different modeling paradigms. The empirical analysis proceeds from traditional linear regression to more sophisticated regularized models (Lasso and Ridge) and machine learning algorithms (Random Forest and Gradient Boosting). The findings reveal several key insights. First, linear models suggest that call and put volumes are consistently significant predictors, while the inclusion of sectoral sentiment—especially EASS indicators—substantially improves explanatory power. Lasso regression confirms these findings and provides a more parsimonious specification, selecting a subset of influential variables that include call, put, RAX, and sector-specific sentiment related to money markets, financials, and cross-asset behavior. Ridge regression complements this by retaining the full set of predictors and offering greater model stability, further confirming the robustness of the identified relationships. Second, while Random Forest models perform poorly in predictive terms—likely due to limited sample size and low signal-to-noise ratio—their variable importance scores align well with linear and penalized models, particularly highlighting the role of call, put, and cross-asset sentiment. Third, Gradient Boosting models achieve nearly perfect in-sample fit ($R^2 \approx 0.999$), identifying a rich structure of nonlinear dependencies between sentiment and returns. Notably, this model reveals that macroeconomic sentiment (RAX, EPU) and equity-specific sentiment (EASSEQUI) are especially influential in explaining large-cap returns (EUROSTOXX50), whereas sectoral sentiment and sovereign risk appear more relevant for the broader market (EUROSTOXX600). Despite these impressive results, the extremely high in-sample performance suggests the need for careful out-of-sample validation to rule out overfitting. Taken together, the evidence supports the conclusion that investor sentiment has significant explanatory power for European equity returns. The consistent relevance of options-based sentiment (call and put volumes), along with sectoral and macroeconomic measures, indicates that behavioral factors contribute meaningfully to return dynamics beyond what can be captured by fundamental or rational expectations models. Moreover, the differential sensitivity of the EUROSTOXX600 and EUROSTOXX50 suggests that market breadth and firm size may condition the strength and nature of sentiment effects.

For investors, these findings offer guidance on incorporating sentiment metrics into forecasting and asset allocation models. For policymakers, the results highlight the potential for real-time sentiment monitoring to serve as an early warning signal of financial stress. Future research should expand this analysis to higher-frequency data, explore alternative sentiment sources (e.g., news analytics, social media), and assess the out-of-sample performance of sentiment-driven strategies.

References

Anastasiou, D., Ballis, A., & Drakos, K. (2022). *Constructing a positive sentiment index for COVID-19: Evidence from G20 stock markets*. International Review of Financial Analysis, 81, 102111.

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

Baker, M., Wurgler, J., Yuan, Y., (2012). *Global, local, and contagious investor sentiment.* Journal of Financial Economics 104, 272–287.

Baker, S.R., Bloom, N., Davis, S.J., (2016). *Measuring Economic Policy Uncertainty.* The Quarterly Journal of Economics, 131 (4), 1593-1636.

Corredor, P., Ferrer, E., & Santamaria, R. (2013). *Investor sentiment effect in stock markets: Stock characteristics or country-specific factors?* International Review of Economics & Finance, 27, 572-591.

Da, Z., Engelberg, J., & Gao, P. (2015). *The sum of all FEARS investor sentiment and asset prices*. The Review of Financial Studies, 28(1), 1-32.

Gao, Z., Ren, H., & Zhang, B. (2020). *Googling investor sentiment around the world*. Journal of Financial and Quantitative Analysis, 55(2), 549-580.

González-Sánchez, M., & Morales de Vega, M. E. (2021). *Influence of Bloomberg's Investor Sentiment Index: Evidence from European Union Financial Sector*. Mathematics, 9(4), 297.

Huang, C., Simpson, S., Ulybina, D., & Roitman, A. (2019). *News-based sentiment indicators*. International Monetary Fund.

Reis, P. M. N., & Pinho, C. (2020). *A new European investor sentiment index (EURsent) and its return and volatility predictability*. Journal of Behavioral and Experimental Finance, 27, 100373.

Sun, L., Najand, M., & Shen, J. (2016). *Stock return predictability and investor sentiment: A high-frequency perspective*. Journal of Banking & Finance, 73, 147-164.

Xu, H. C., & Zhou, W. X. (2018). *A weekly sentiment index and the cross-section of stock returns*. Finance Research Letters, 27, 135-139.