International Review of Financial Analysis 94 (2024) 103315

Contents lists available at ScienceDirect

International Review of Financial Analysis

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

Analyzing credit spread changes using explainable artificial intelligence

Julia Heger [a,1], Aleksey Min [b,*,1], Rudi Zagst [b,1]

[a] University of Augsburg, Chair of Analytics & Optimization, Universitätsstraße 16, Augsburg, 86159, Germany
[b] Technical University of Munich, Chair of Mathematical Finance, Parkring 11, Garching-Hochbrück, 85748, Germany

ARTICLE INFO

JEL classification:
C10
C18
C45
C52
C53
C58

Keywords:
Credit spread changes
Random forest
Partial dependence plot
H-statistic
SHAP values
Hedging

ABSTRACT

We compare linear regression, local polynomial regression and selected machine learning methods for modeling credit spread changes. Using partial dependence plots (PDPs) and H-statistic, we find that the outperformance of machine learning models compared to regression ones is mostly attributable to complex non-linearities and not to interactions. The PDPs are additionally used to perform a factor hedging. For the first time, credit spread changes are decomposed by applying SHapley Additive exPlanation (SHAP) values. The proposed framework is applied to US and Euro Area corporate and covered bond credit spread changes of different maturities to quantify the influence of several macroeconomic and financial variables. Despite several commonalities between the decompositions of US and Euro Area credit spread changes, we also observe some differences — particularly related to the impact of certain explanatory variables during crisis periods.

1. Introduction

Credit spreads are defined as the difference between interests paid on bonds with a low and a higher level of risks. Fluctuations in credit spreads indicate that general market conditions are changing. These fluctuations appear frequently due to various reasons such as changes in economic conditions, changes in liquidity etc. The seminal work by Collin-Dufresne et al. (2001) shows that factors, which in theory should explain credit spread changes, fail to do so since they only account for about one quarter of the variation in credit spread changes according to the adjusted R^2. Moreover, these authors find that regression residuals are mostly driven by a single common factor and suppose that a measure of supply/demand shocks can be a pertinent determinant of credit spread changes. Since the revelation of these surprising results, researchers are eager to find the true unknown common factor.

Among them, Javadi et al. (2018) study the impact of Federal Open Market Committee (FOMC) actions as potential determinants of changes in credit spreads and lend credence to the conjecture of Collin-Dufresne et al. (2001). Namely, they identify FOMC actions, which by design affect the supply/demand of bonds, as determinants of credit spread changes. However, when adding FOMC action dummies, the whole set of explanatory variables still only accounts for at most 22% of the variation in credit spread changes according to the adjusted R^2 for the data set considered by the authors.

Friewald and Nagler (2019) examine, whether systematic over-the-counter market frictions drive the unexplained common component proven to be present in yield spread changes. They claim that market-wide inventory, search, and bargaining frictions explain about one third of the variation of the common component. Yet, when additionally adding all of these frictions, the whole set of factors still only accounts for about 30% of the variation in credit spread changes according to the adjusted R^2.

Kaviani et al. (2020) study the impact of policy uncertainty on corporate credit spread changes. The policy uncertainty is measured by a policy uncertainty index (PUI), that is constructed from four components, namely a comprehensive measure of uncertainty based on the number of articles about economic policy uncertainty in ten major newspapers, taxation uncertainty based on data from the Congressional Budget Office on expiring tax provisions, inflation uncertainty, and government expenditure uncertainty based on the dispersion in professional analyst forecasts. The authors find a significantly positive relationship between the PUI and changes in corporate credit spreads. Adding the PUI as variable to a regression including the control

* Corresponding author.
E-mail addresses: julia.heger@uni-a.de (J. Heger), min@tum.de (A. Min), zagst@tum.de (R. Zagst).
[1] The author has nothing to disclose.

https://doi.org/10.1016/j.irfa.2024.103315
Received 13 June 2023; Received in revised form 31 October 2023; Accepted 12 April 2024
Available online 16 April 2024

variables considered in Collin-Dufresne et al. (2001) yields an adjusted R^2 of at most 31%.

Likewise, He et al. (2022) investigate whether two intermediary-based factors, a broad financial distress measure and a dealer corporate bond inventory measure are responsible for the unknown common component in credit spread changes. They find that those two factors explain about half of the unexplained common variation in credit spread changes. When adding these factors, the whole set of explanatory variables accounts for about 50% of the variation in credit spread changes according to the adjusted R^2.

In addition, besides the application of regression modeling, the literature offers further approaches used for analyzing credit spread changes. For example, Manzoni (2002) follow a time-series approach by applying ARCH and GARCH models to changes in the credit spreads on the Sterling Eurobond index. The authors find that the time-series properties of the credit spreads provide evidence of non-linearities.

Despite multiple attempts examining a wide range of possible explanatory variables, the large common component remains unexplained. In this paper, we examine, whether the lack in performance when modeling credit spread changes using linear regression on the standard set of explanatory variables is due to non-linearities and variable interactions and provide strong empirical evidence that this is indeed the case. We show how machine learning (ML) models can be applied to decompose credit spread changes and propose alternative approaches for performing a factor hedging of credit spread changes.

To be precise, our contributions are fourfold. Firstly, we link the area of analyzing the determinants of credit spread changes with the field of machine learning. We conduct a comparative analysis of a selection of the most common machine learning methods. The ML models are additionally compared to a standard linear regression model and a local polynomial regression, the first serving as benchmark. We identify the best-performing method using performance measures such as the coefficient of determination R^2, the Sum of Squares Residual (SSR) ratio, the root-mean-squared error (RMSE) and the direction-of-change (DoC) accuracy. The performance is evaluated using 20 data sets containing credit spread change data of different bond types, maturities and countries, each covering a period of twenty years.

Second, we apply two explainable artificial intelligence (XAI) techniques called partial dependence plots (PDPs) and H-Statistic — introduced by Friedman (2001) and Friedman and Popescu (2008), respectively — to examine if the outstanding performance of the machine learning models is attributable to non-linear relationships or interactions. We additionally make use of the PDPs to provide evidence that our approach can be used for hedging.

Third, we propose a novel approach for quantifying the influence of different macroeconomic and financial variables on credit spread changes. Our selected variables are similar to those of Collin-Dufresne et al. (2001). Unlike the techniques of all other studies known to us, our method is based on an XAI technique called SHapley Additive exPlanation (SHAP) values. In contrast to the classic regression-based decomposition approach pursued by Collin-Dufresne et al. (2001), which provides one decomposition of the credit spread changes over the whole time period, our approach enables the explanation of every single credit spread change. Moreover, as SHAP values are model-agnostic, our approach is not limited to linear regression models but also allows for the usage of machine learning models.

Fourth, the above introduced methodology of modeling the credit spread changes using a supervised machine learning method, examining the presence of non-linearities and interactions, executing a factor hedging and decomposing the model predictions into the spread changes' determinants is also validated on US and Euro Area corporate and covered bond spread changes of different maturities.

So overall, the paper sheds light on the questions, which method to use when it comes to modeling credit spread changes and why this method is best suited. Besides, it offers an XAI method that can be used for executing a factor hedging. Finally, it also proposes an approach for decomposing credit spread changes into its components. This approach allows for the usage of ML models and additionally enables a decomposition of every single credit spread change and not only of one decomposition over the whole time period considered in the model.

The remainder of this paper proceeds as follows: In Section 2, we briefly present the machine learning methods used for modeling the credit spread changes, followed by a short summary of the theory required for the analysis of non-linearities and interactions as well as for the factor hedging. The section ends with a concise description of the theory used for the decomposition methodology. In Section 3, the previously introduced methods are applied to model credit spread changes, examine the presence of non-linearities and interactions, perform a factor hedging and decompose the model predictions into the spread changes' determinants. To be precise, this is done for corporate and covered bond spread changes across different maturities and countries. Section 4 concludes and discusses further research.

2. Modeling methodology

We start by outlining the machine learning methods considered in our analysis. Some of them are standard statistical models, whereas others are more advanced ML methods. Subsequently, the performance measures used to select the best model are presented. We statistically justify the choice of the best model using Friedman's rank sum test with Iman and Davenport correction followed by Friedman's post-hoc test. Afterwards, we introduce our framework for validating the presence of non-linearities and interactions. Finally, we illustrate the methodology used to decompose credit spread changes.

2.1. Machine learning models

In order to decompose credit spread changes using machine learning methodology, clearly the first step is to find the model that shows the best possible performance in modeling the credit spread changes. This is done by conducting a comparative study of the five popular machine learning methods listed in Table 1.

Most of the studies analyzing credit spread determinants typically base their analysis on the application of linear regression models, see for example Collin-Dufresne et al. (2001), Friewald and Nagler (2019) or He et al. (2022). Therefore, we additionally fit a standard linear regression, which is considered as benchmark model and a local polynomial regression. In particular, we check whether and to which extend advanced machine learning models are able to outperform the classic OLS regression and the local polynomial regression in modeling credit spread changes. As a first machine learning model we include a common support vector regression in our analysis. In addition, as suggested by some literature on modeling bond returns, credit spreads or credit default swaps using machine learning such as Bianchi et al. (2021), Xiong et al. (2019) or Son et al. (2016), respectively, we consider two different ensemble methods — random forest and Bayesian additive regression trees. Besides these ensemble methods, the literature also suggests using neural networks. Therefore, we additionally consider a Bayesian regularized artificial neural network as well as a long short-term memory neural network. A brief description of these models is given in Table 1.

2.2. Performance evaluation

First, we evaluate the models predictive ability in terms of the root-mean-squared error. Second, to ensure that our model correctly captures the direction of the spread changes, we consider the direction-of-change accuracy by binary labeling the series according to whether its values are positive or negative and evaluating the corresponding accuracy. Third, to capture the proportion of variance in credit spread changes that is explained by the explanatory variables, we evaluate the

Table 1
Overview of the considered ML models.

Model	Description
OLS	OLS refers to the classic multiple linear regression, see for example Chapter 3 in Greene (2003).
LPR	Local Polynomial Regression (LPR) is a non-parametric regression method where a weighted polynomial regression is fitted at each point of the independent variable based on all the observations, whose values of the independent variable are within the range of a predefined bandwidth, see for example Cleveland et al. (1992).
SVR	Support vector regression (SVR) is the regression version of a support vector machine (SVM), which belongs to the so-called large margin classifiers. The SVM aims at separating instances of different classes by the maximum margin hyperplane, where the margin is the minimum distance from instances of different classes to the classification hyperplane. In contrast, SVR aims at finding a hyperplane that is closest to all data points with respect to a chosen distance measure. For more details, we refer to Section 14.5 in Murphy (2012).
RF	The random forest (RF) belongs to the class of bagging methods. It bootstraps the underlying sample and estimates one model — typically a decision tree — for each bootstrap sample by using only a small, randomly selected subset of the available features. The final predictions are obtained by averaging all models' predictions. For more information, we refer to Chapter 3 of Zhou (2012).
BART	The Bayesian additive regression tree model (BART) is defined by a sum-of-trees model paired with a regularization prior on the model parameters. Thereby, each tree is constrained by a regularization parameter to keep the individual effect of each tree small. For more details on BART models, we refer to Chipman et al. (2010).
BRANN	The Bayesian regularization artificial neural network (BRANN) consists of a common artificial neural network (NN) paired with Bayesian regularization. Thereby, one additionally imposes a certain prior distribution on the NN model parameters and penalizes large weights. The advantage of such BRANNs is their robustness. For more information on BRANNs, see for example Sariev and Germano (2019).
LSTM	The key idea behind long short-term memory (LSTM) NN is the so-called cell state that runs through the whole NN with only linear interactions. Thereby, it allows the information to pass straight through the units and serves as a highway for the gradients to flow during backpropagation. These linear interactions are regulated by three gates — forget gate, input gate, and output gate — that exhibit the ability to add or remove information to the cell state. For more details on LSTMs, see for example Hochreiter and Schmidhuber (1997) or Gers et al. (2000).

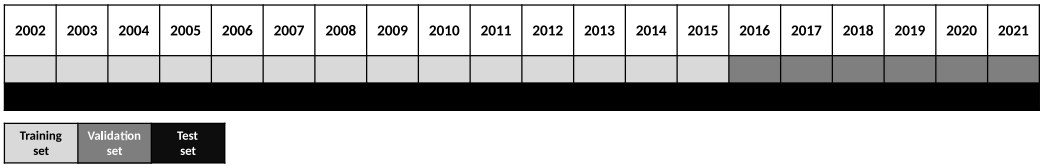

Fig. 1. Splitting data for in-sample performance evaluation into training, validation and test sets using an extending window approach.

coefficient of determination R^2 in terms of in-sample performance. In contrast to the in-sample R^2, the out-of-sample R^2 considered by Campbell and Thompson (2008) might take values in $(-\infty, 1]$ and is not simple to interpret. However, note that comparing two models m_1 and m_2 with respect to the out-of-sample R^2 is equivalent to comparing their Sum of Squares Residuals:

$$R^2_{m_1} < R^2_{m_2} \Leftrightarrow 1 - \frac{SSR_{m_1}}{SST} < 1 - \frac{SSR_{m_2}}{SST} \Leftrightarrow SSR_{m_2} < SSR_{m_1},$$

where SST denotes the Sum of Squares Total. In order to enhance interpretability, we propose to examine the SSR ratio SSR_{m_2}/SSR_{m_1} with m_1 being the linear regression model. This quantity indicates the percentage of the SSR of model m_2 compared to the SSR of the linear regression model and thereby also allows for the comparison of two different models other than OLS via the linear regression.

To evaluate the in-sample performance, we first need to obtain suitable hyperparameters for each of the considered ML models listed in Table 1. To do so, we split the whole data into training and validation set, successively fit the models for different hyperparameters on the training set as illustrated in Fig. 1 and predict the validation set month-by-month whereby the training set is enlarged by one observation every month in an extending window manner. Next, the hyperparameters that perform best over this six-year validation period are chosen for each model. Finally, the in-sample performance is obtained by fitting the model using the chosen hyperparameters and evaluating the performance metrics each over the whole data set.

To determine the out-of-sample performance, we adapt the idea introduced in Gu et al. (2020). Doing so, each data set is divided into training, validation and test sets in an extending window approach as visualized in Fig. 2. The first training set comprises 70% of the data and subsequently enlarges step by the data of the upcoming year. In contrast, the validation set always comprises four years and the test set consists of the year following the validation set whereby both, validation and test set move onward in time. In every step, i.e. every row of Fig. 2, all models are fitted on the training set for a variety of different hyperparameter combinations. The validation set is predicted month-by-month. This is done by refitting the models every month on the monthly enlarging training set and predicting the month following the current training set. This procedure is repeated over the whole four-year validation period. Subsequently, the performance is evaluated by comparing the series of predictions with their corresponding actuals. Then, the hyperparameters exhibiting the best performance per model are chosen and used to predict the test-year month-by-month. The final out-of-sample performance is obtained by evaluating the performance measures on the test period.

In order to validate whether our selected best model is indeed *significantly* better than the other considered models in terms of R^2 respectively SSR ratio, RMSE and DoC accuracy, we apply Friedman's rank sum test with Iman and Davenport correction followed by Friedman's post-hoc test with a control model as proposed by Demsar (2006) [Section 3.2.2].

If the null-hypothesis that all algorithms are equivalent can be rejected, one can proceed with Friedman's post-hoc test with control classifier in order to confirm that the control model significantly outperforms all other models, see Demsar (2006) [Section 3.2.2]. To account for multiple comparisons, finally Benjamini–Hochberg's multiple testing correction proposed by Benjamini and Hochberg (1995) is used.

2.3. Methodology for examining the presence of non-linearities and interactions

In this section, we present our methodology for examining the presence of non-linearities and interactions in the model using two explainable AI techniques called partial dependence plot and H-statistic. We additionally shed light on the usage of partial dependence plots to perform a factor hedging.

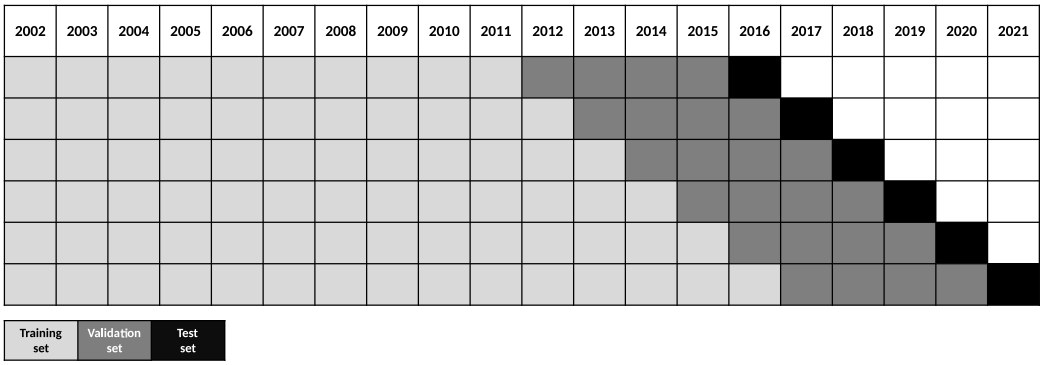

Fig. 2. Splitting data for out-of-sample performance evaluation into training, validation and test sets using an extending window approach.

The partial dependence plot introduced by Friedman (2001) depicts the marginal effect of an explanatory variable on the predicted credit spread changes and thereby allows to examine the functional relationship between the considered explanatory variable and the change in credit spreads. That is to say, when applied to a linear regression model the partial dependence profile shows a linear relationship whereas this might not be the case when applied to machine learning models that are able to learn non-linearities.

Formally, let x_S be the explanatory variable for which the partial dependence should be evaluated and let x_C be all other explanatory variables. Then, the partial dependence function of feature x_S for model $\hat{f}$ is defined as

$$\hat{f}_S(x_S) = E_{x_C}[\hat{f}(x_S, x_C)] = \int \hat{f}(x_S, x_C) p_C(x_C) dx_C. \tag{1}$$

In the above formula, $p_C(x_C) = \int p(x_S, x_C) dx_S$ is the marginal probability density of x_C with $p(x_S, x_C)$ representing the joint density of all inputs. The marginal probability density can be estimated on the training data. Then, Eq. (1) becomes

$$\hat{f}_S(x_S) = \frac{1}{n} \sum_{i=1}^{n} \hat{f}(x_S, x_{C_i}),$$

with n being the number of observations. This partial dependence function evaluates the average marginal effect on the prediction of the credit spread change for given values of explanatory variable x_S. Intuitively, the partial dependence function at a certain value of the explanatory variable is calculated as the average prediction when all data points are forced to assume this value of the explanatory variable. However, the PDP assumes that explanatory variables x_S for which the partial dependence is evaluated are uncorrelated with all other features x_C. Otherwise, the partial dependence plot might include unlikely data points. For more information, we refer the interested reader to Molnar (2022).

To ensure that the assumption of uncorrelated features is satisfied, only explanatory variables with rather small correlations should be chosen. Table 4 shows two correlation matrices whose cells represent the pairwise correlations between all explanatory variables used in our models during the case study for the US and the Euro Area case, respectively. As can be seen, overall the correlations between the considered features are quite small — for example for the US data, the correlation between ΔCLI and ΔBidAsk only amounts to 0.003. Hence, the assumption of uncorrelated features can be assumed to approximately hold for our case studies.

Besides examining the functional form of the relationship between bond spread changes and explanatory variables, PDPs can also be used to perform a factor hedging using credit default swaps (CDS). To see this, note that a PDP can be interpreted as the sensitivity of the change in the credit spread depending on the change of some explanatory variable and thus represents a hedge ratio.

To execute a hedging, a model for the credit default swap spread changes must first be fitted using the explanatory variables of the credit spread change model. In the linear regression framework, this approach is pursued to select explaining factors for changes in credit default swap spreads among macroeconomic variables considered for credit spread changes (see e.g. Breitenfellner & Wagner, 2012). Hence, it is reasonable to apply the same type of model with the same explanatory variables on credit default swap spread changes that we used for the credit spread changes.

After fitting the model, the PDPs for the credit default swap spread changes must be created for all considered explanatory variables similar to the PDPs for the credit spread changes. These credit spread change and credit default swap spread change PDPs then provide the sensitivities $PDP^{CS}_{x_i} = \frac{\Delta CS}{\Delta x_i}$ and $PDP^{CDS}_{x_i} = \frac{\Delta CDS}{\Delta x_i}$ relative to the explanatory variables x_i, respectively.

Having these sensitivities available, the goal of a *factor hedging* is to determine the factor α_i such that the squared error between the credit spread PDP and the credit default swap spread PDP is minimized, i.e., more formally:

$$\min_{\alpha_i} \sum_{k \in K} (PDP^{CS}_{x_{ik}} - \alpha_i \cdot PDP^{CDS}_{x_{ik}})^2,$$

where K represents the set of data points of the PDPs whose x-values are contained within the range between their 5th and 95th quantile. The factor α_i can be determined by fitting a linear regression with target $PDP^{CS}_{x_i}$ and explanatory variable $PDP^{CDS}_{x_i}$. To be precise, α_i is then given by the slope of the linear regression and implies that α_i parts of the credit default swap must be purchased for hedging.

Note that such a hedging can be transformed into a cash payment. However, this would require some information regarding the composition of the bond portfolio. Since the paper at hand uses bond indices, this required information is not available in our case. Hence, this paper solely focuses on the hedging execution using PDPs.

The H-statistic introduced by Friedman and Popescu (2008) is based on the idea of partial dependence and allows to measure whether and to what extend an explanatory variable interacts with all other variables. Formally, the H-statistic for testing whether one variable x_S interacts with any other variable is defined as

$$H^2_S = \frac{\sum_{i=1}^{n} [\hat{f}(x_{S_i}, x_{C_i}) - \hat{f}_S(x_{S_i}) - \hat{f}_C(x_{C_i})]^2}{\sum_{i=1}^{n} \hat{f}^2(x_{S_i}, x_{C_i})},$$

where $\hat{f}(x_{S_i}, x_{C_i})$ represents the model prediction for the whole original observation i and $\hat{f}_S$ and $\hat{f}_C$ denote the partial dependence functions for feature S and feature subset C respectively. The H-statistic differs from zero to the extent that the explanatory variable x_S interacts with one or more other features, see Molnar (2022). As the H-statistic is based on partial dependence, it overtakes the assumption of uncorrelated features. As mentioned before, we account for this requirement by choosing a set of explanatory variables with rather small correlations as can be seen in Table 4.

2.4. Methodology for decomposing credit spread changes

In this section, we present our methodology for decomposing credit spread changes that are modeled using advanced ML methods. These methods are known to be black boxes, as — in contrast to linear models — it is difficult to interpret them directly due to their high complexity. Fortunately, in the last few years, a variety of explainable artificial intelligence methods for interpreting black-box models evolved. Among them, SHAP values that are proposed by Lundberg and Lee (2017) are current state-of-the art in the field of model interpretability.

SHAP values approximate the model locally by using a so-called explanation model g. To be precise, the SHAP values belong to the class of additive feature attribution methods and aim at explaining the original model f locally around the instance x under consideration by using an explanation model g, that is a linear function of binary variables

$$g(x') = \phi_0 + \sum_{i=1}^{|F|} \phi_i(x) x'_i, \tag{2}$$

where the base value $\phi_0 := E[f(x)]$ is defined as the average model prediction, $\phi_i \in \mathbb{R}$ represents the attribution of feature i and $x' \in \{0,1\}^{|F|}$ represents the so-called simplified inputs with $|F|$ indicating the number of such simplified features, which are defined to be binary with an entry of 1 implying the presence of the corresponding feature value, whereas an entry of 0 represents its absence. This local explanation model satisfies three desirable properties, namely local accuracy, missingness and consistency.

The **local accuracy** requires the explanation model g to at least match the output of the original model f for the simplified input x', i.e.

$$f(x) = g(x') = \phi_0 + \sum_{i=1}^{|F|} \phi_i(x) x'_i \tag{3}$$

when $x = h_x(x')$. The mapping h_x transfers the simplified input x' into the original feature space as detailed in the below description of the KernelSHAP.

The **missingness** property requires missing features to have zero impact, i.e.

$$x'_i = 0 \Longrightarrow \phi_i(x) = 0. \tag{4}$$

The **consistency** states that if a model changes in a way such that some simplified input's contribution increases or stays the same regardless of the other inputs, then that input's attribution should not decrease, i.e. for any two models $\hat{f}$ and f,

$$\hat{f}_S - \hat{f}_{S-i} \geq f_S - f_{S-i} \Longrightarrow \hat{\phi}_i(x) \geq \phi_i(x), \tag{5}$$

where f_S denotes f being applied to features contained in $S \subseteq F$.

According to Lundberg and Lee (2017), the unique explanation model belonging to the class of additive feature attribution methods and satisfying the three properties mentioned above is given by

$$\phi_i(x) = \sum_{S \subseteq F \setminus \{i\}} \frac{|S|!(|F| - |S| - 1)!}{|F|!} [f(x_S) - f(x_{S-i})], \tag{6}$$

where $x_S := (x_{k(1)}, \ldots, x_{k(|S|)}), 1 \leq k(1) \leq k(|S|) \leq |F|$ denotes the observation x under consideration where only features in S are considered and $f(x_S) := E[f(x)|x_S]$ denotes the conditional expectation given x_S.

The goal of SHAP values is to attribute a value to each explanatory variable that corresponds to the change in the expected model prediction. To be precise, SHAP values explain the difference between the base value, that is to say the average model prediction $E[f(x)]$ and the current output $f(x)$. This is qualitatively visualized in Fig. 3 with $f(x)$ being the modeled spread change and $E[f(x)]$ being its average.

Changes in credit spreads are calculated by taking the first order differences between consecutive observations. The resulting credit spread changes are stationary with zero mean. Therefore, SHAP values indeed explain the difference between the average modeled spread change and each individual modeled spread change.

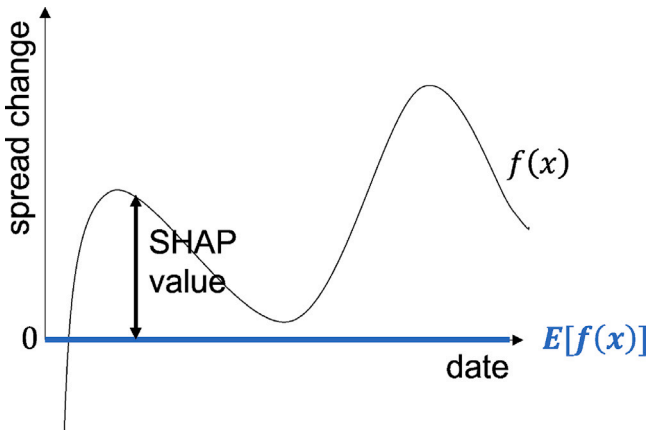

Fig. 3. Decomposition of credit spread changes using SHAP values.

As the exact calculation of SHAP values is challenging, Lundberg and Lee (2017) propose a model-agnostic approach called KernelSHAP and a variety of model-specific approaches such as DeepSHAP for deep learning models to approximate SHAP values. Moreover, Lundberg et al. (2018) propose the model-specific TreeSHAP, a fast implementation for tree-based models. As according to our study, the best method for modeling credit spread changes is the random forest, which is a tree-based method, both the KernelSHAP and the TreeSHAP can be applied for the decomposition of the spread changes. To decide which method to use, we compare the advantages and disadvantages of both approaches.

The biggest advantage of SHAP values in general is their strong game-theoretical foundation ensuring that the prediction is fairly distributed across all explanatory variables. However, SHAP values are computationally expensive. Particularly the KernelSHAP is pretty slow making it difficult to apply this method to a large number of instances. Moreover, in case of feature dependence the random replacement of feature values carried out in the KernelSHAP might result in putting weight on unlikely observations. The TreeSHAP overcomes this issue by modeling the conditional expected prediction. However, by relying on conditional expected predictions the TreeSHAP changes the value function, which might result in unintuitive attribution scores including non-zero scores for features actually having no influence at all. For a more detailed comparison of advantages and disadvantages of SHAP values see Molnar (2022).

In our study, we only execute the decomposition of the credit spread changes once per target. Moreover, our data set only comprises 240 observations per target. Thus, the computational complexity does not matter in our case. In addition, as our goal is to obtain a reliable decomposition of the credit spread changes, we do not want to get unintuitive attribution scores and therefore rather rely on the KernelSHAP method. To overcome the issue of the KernelSHAP assuming feature independence, we choose a selection of explanatory variables that are rather low correlated as can be seen in Table 4.

Intuitively, the KernelSHAP proceeds as depicted in Fig. 4. Before the KernelSHAP procedure can be applied, one needs to choose the observation that should be decomposed. Moreover, the training and background sets need to be defined whereby the background set either equals the full training set or a random subset thereof. Lastly, the ML model that should be explained must be trained on the training set. After having fulfilled these requirements, the KernelSHAP procedure can be applied. Doing so, the first step is to construct possible coalition vectors $x'_k \in \{0,1\}^{|F|}$, where 1 again indicates the presence of a feature and 0 its absence. Next, the coalition vectors need to be transferred into the original feature space. This is done by replacing each 1 with

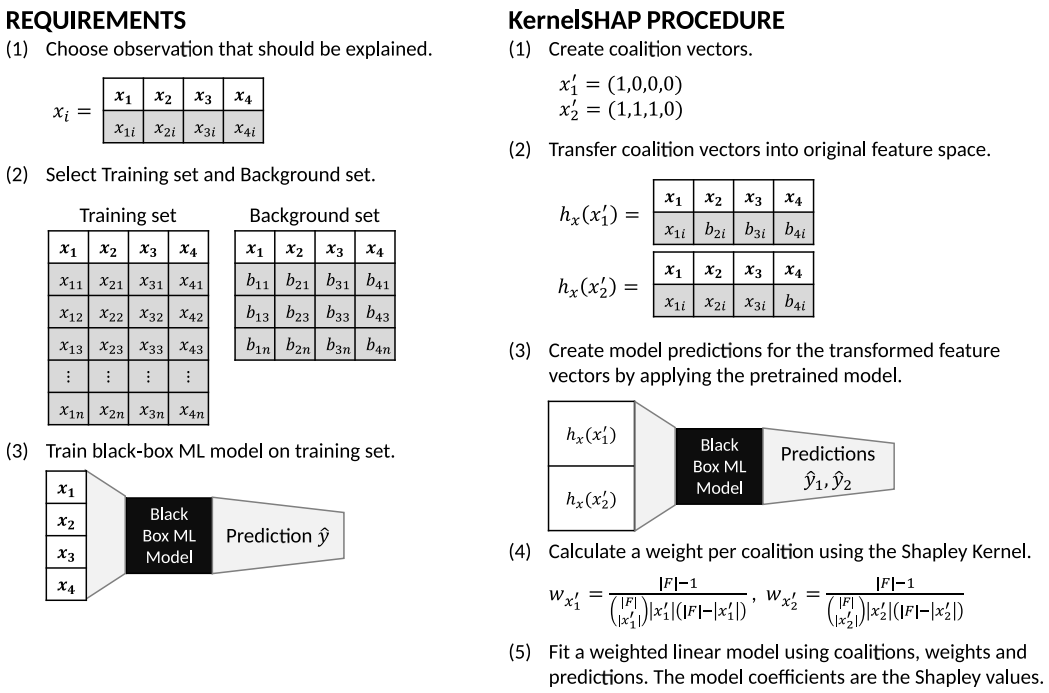

Fig. 4. KernelSHAP procedure.

the original feature value of the observation that should be explained and each 0 by a random feature sample from the background set. Subsequently, the pretrained ML model is applied to the backtransformed features and a prediction is generated. Next, using the Shapley Kernel a weight is calculated for each coalition. Intuitively, the Shapley Kernel ensures that those coalitions where either almost all or almost no features are present get the highest weights. Finally, the coalitions, weights and predictions are used to fit a weighted linear model. The resulting model coefficients are the Shapley values.

3. Application

In the following, the methodology introduced in Section 2 is applied to US and Euro Area corporate and covered bond credit spread changes of different maturities.

3.1. Data

Recall that the credit spread is defined as the difference between a corporate respectively covered bond yield and the corresponding sovereign bond yield of the same maturity. We use monthly IBoxx Euro and USD corporate, covered and sovereign data for different maturities that is provided by Datastream. Based on this, we define the credit spread change ΔCS_t at time t as the difference between two consecutive credit spread observations, i.e.

$$\Delta CS_t = CS_t - CS_{t-1}.$$

The monthly data covers the range from January 2002 until December 2021. An overview of the considered target variables is provided in Table 2.

To fit the previously introduced models on credit spread changes, we do not only use historical spread changes but supplement the above mentioned data by a set of explanatory variables provided by Datastream. The selection of explanatory variables is mostly composed of the variables included in the benchmark linear regression executed by Collin-Dufresne et al. (2001), Friewald and Nagler (2019) and He et al. (2022). These variables include changes in the 10-year Benchmark Treasury rate, squared changes in the 10-year Benchmark Treasury rate, changes in the slope of the yield curve defined as the difference between 10-year and 2-year Benchmark Treasury yields, changes in VIX, S&P 500 returns and changes in the SKEW index. The latter, in line with Kim et al. (2017), is a proxy for the slope of the volatility smirk.

Moreover, as Martell (2008) finds evidence of a liquidity-based explanation for the unknown common component in credit spread changes, we augment the explanatory variables used in the above mentioned literature by the bid–ask spread as proxy for bond market liquidity as proposed by Amihud and Mendelson (1986) and Feldhütter and Poulsen (2018). In addition, we add OECD's Composite Leading Indicator (CLI) to capture early signals of turning points in the business cycle as suggested by Hauptmann et al. (2014). An overview of the considered explanatory variables is given in Table 3; their correlations are summarized in Table 4. This data is used to fit the previously mentioned models listed in Table 1.

3.2. Model performance comparison

The models' in-sample as well as out-of-sample performances in terms of the coefficient of determination R^2, the SSR ratio, the root-mean-squared error and the direction-of-change accuracy are evaluated across all considered data sets. The in-sample performance is obtained by fitting the models and evaluating the above mentioned metrics each on the whole 20-year time period from January 2002 to December 2021 whereas the out-of-sample performance is evaluated using the extending window approach introduced in Section 2.2. The resulting performances with respect to the different measures are summarized in Tables 5–7.

Taking a look at the R^2 listed in Table 5, the by far best model in terms of the coefficient of determination is the random forest model showing an in-sample R^2 amounting to around 0.773 across all considered data sets. This implies that the proportion of variance in the credit spread changes that is explained by the explanatory variables introduced in Section 3.1 constitutes around 77% for the random forest model, while there is only a mere 23% unexplained variation. Moreover, notice that almost all machine learning models at least double performance in terms of in-sample R^2 compared to the linear regression

Table 2
Overview of monthly target variables.

Name	Description
$\Delta CS^{\text{USD Corporate 1-3}}$	USD corporate bond spread changes (1–3 years)
$\Delta CS^{\text{USD Corporate 3-5}}$	USD corporate bond spread changes (3–5 years)
$\Delta CS^{\text{USD Corporate 5-7}}$	USD corporate bond spread changes (5–7 years)
$\Delta CS^{\text{USD Corporate 7-10}}$	USD corporate bond spread changes (7–10 years)
$\Delta CS^{\text{USD Corporate 10+}}$	USD corporate bond spread changes (10+ years)
$\Delta CS^{\text{USD Covered 1-3}}$	USD covered bond spread changes (1–3 years)
$\Delta CS^{\text{USD Covered 3-5}}$	USD covered bond spread changes (3–5 years)
$\Delta CS^{\text{USD Covered 5-7}}$	USD covered bond spread changes (5–7 years)
$\Delta CS^{\text{USD Covered 7-10}}$	USD covered bond spread changes (7–10 years)
$\Delta CS^{\text{USD Covered 10+}}$	USD covered bond spread changes (10+ years)
$\Delta CS^{\text{EURO Corporate 1-3}}$	Euro corporate bond spread changes (1–3 years)
$\Delta CS^{\text{EURO Corporate 3-5}}$	Euro corporate bond spread changes (3–5 years)
$\Delta CS^{\text{EURO Corporate 5-7}}$	Euro corporate bond spread changes (5–7 years)
$\Delta CS^{\text{EURO Corporate 7-10}}$	Euro corporate bond spread changes (7–10 years)
$\Delta CS^{\text{EURO Corporate 10+}}$	Euro corporate bond spread changes (10+ years)
$\Delta CS^{\text{EURO Covered 1-3}}$	Euro covered bond spread changes (1–3 years)
$\Delta CS^{\text{EURO Covered 3-5}}$	Euro covered bond spread changes (3–5 years)
$\Delta CS^{\text{EURO Covered 5-7}}$	Euro covered bond spread changes (5–7 years)
$\Delta CS^{\text{EURO Covered 7-10}}$	Euro covered bond spread changes (7–10 years)
$\Delta CS^{\text{EURO Covered 10+}}$	Euro covered bond spread changes (10+ years)

Table 3
Overview of monthly explanatory variables.

Name	Description
Δr^{10}	change in 10-year Benchmark Treasury rate
$(\Delta r^{10})^2$	squared change in 10-year Benchmark Treasury rate
ΔSlope	change in 10-year minus 2-year Treasury yields
ΔVIX	change in volatility index VIX
$\text{ret}^{\text{SP500}}$	S&P 500 returns
ΔSKEW	change in SKEW
ΔBidAsk	change in Bid–Ask Spread
ΔCLI	change in OECD Composite Leading Indicator

Table 4
Correlation matrix.

US								
	Δr^{10}	$(\Delta r^{10})^2$	ΔSlope	ΔVIX	$\text{ret}^{\text{SP500}}$	ΔSKEW	ΔBidAsk	ΔCLI
Δr^{10}	1							
$(\Delta r^{10})^2$	−0.065	1						
ΔSlope	0.396	0.049	1					
ΔVIX	−0.152	0.131	0.015	1				
$\text{ret}^{\text{SP500}}$	0.213	−0.137	0.058	−0.541	1			
ΔSKEW	0.119	−0.137	0.109	−0.157	0.235	1		
ΔBidAsk	0.010	−0.115	0.029	−0.009	0.000	0.022	1	
ΔCLI	0.159	−0.076	0.086	−0.088	0.230	0.063	0.003	1
Euro Area								
	Δr^{10}	$(\Delta r^{10})^2$	ΔSlope	ΔVIX	$\text{ret}^{\text{SP500}}$	ΔSKEW	ΔBidAsk	ΔCLI
Δr^{10}	1							
$(\Delta r^{10})^2$	−0.128	1						
ΔSlope	0.338	0.062	1					
ΔVIX	−0.175	0.126	0.060	1				
$\text{ret}^{\text{SP500}}$	0.220	−0.181	0.004	−0.541	1			
ΔSKEW	0.142	−0.129	0.114	−0.157	0.235	1		
ΔBidAsk	0.005	0.020	−0.011	−0.092	0.068	−0.017	1	
ΔCLI	0.178	−0.123	0.036	−0.088	0.230	0.063	0.032	1

model, whose in-sample R^2 on average amounts to approximately 18%. This amount of variance explained by the linear regression is in line with the findings of Collin-Dufresne et al. (2001), Friewald and Nagler (2019) and He et al. (2022).

Notably, the random forest model also exhibits a comparably good performance in terms of average out-of-sample SSR ratio, which totals 0.512 implying that the out-of-sample SSR of the random forest model only amounts to approximately 50% of the SSR of the linear regression model. The same holds for the LPR, the LSTM and the SVR model whereas the BART and the BRANN models both perform worse with SSR ratios of 0.629 and 0.911, respectively.

Examining the in-sample as well as out-of-sample RMSE listed in Table 6, the worst model in terms of in-sample RMSE is the local polynomial regression with an average RMSE amounting to 0.176 whereas the worst model in terms of out-of-sample RMSE is the linear regression, showing an average RMSE of 0.199. The by far best-performing model in terms of in-sample RMSE is the random forest model, whose average in-sample RMSE amounts to around 0.090. The best models in terms of out-of-sample RMSE are the local polynomial regression and the random forest with average out-of-sample RMSE of 0.132 and 0.133, respectively.

As can be seen in Table 7, the random forest model also shows the best in-sample as well as out-of-sample performance in terms of direction-of-change accuracy, which on average amounts to around 88.7% respectively 56.6% across all considered data sets. In contrast, the worst model in terms of in-sample and out-of-sample direction-of-change accuracy is the linear regression with average accuracy scores amounting to 58.6% and 52.8%, respectively.

In total, considering all of the above performance measures, the random forest model on average shows the highest R^2, one of the highest SSR ratios, one of the lowest RMSEs and the highest direction-of-change accuracy. Nevertheless, we shall additionally check whether the random forest model is *significantly* better than the other considered models in terms of R^2, SSR ratio, RMSE and direction-of-change accuracy.

To do so, we first use Friedman's rank sum test with Iman and Davenport correction to check the null hypothesis that all considered models are equivalent. Applying this test to the above R^2, RMSE and direction-of-change accuracy in-sample performance tables yields that the null hypothesis can be rejected in all three cases with p-values less than $2.2 \cdot 10^{-16}$. Likewise, applying the test to the SSR ratio, RMSE and direction-of-change accuracy out-of-sample performance tables yields that the null hypothesis can be rejected in case of RMSE and SSR ratio with p-values less than $2.2 \cdot 10^{-16}$. The p-value in case of the direction-of-change accuracy amounts to 0.144.

Therefore, in order to validate whether the random forest model significantly outperforms all other models, we can now proceed with Friedman's post-hoc test using random forest as control model. Applying this test to the R^2, RMSE and direction-of-change accuracy in-sample performance tables yields, that the random forest model is significantly better than all other models with respect to all three performance measures at a ten percent level of significance. Likewise, applying this test to the above SSR ratio, RMSE and direction-of-change accuracy out-of-sample performance tables yields that the random forest model is significantly better than all other analyzed models except for LPR, SVR and LSTM with respect to at least two of the three performance measures at a ten percent level of significance. Moreover, the random forest model shows the best average in-sample performance for all three performance measures and the best out-of-sample performance for one out of the three performance measures. Hence, all in all, the hypothesis tests confirm that the random forest model is better than all other considered models. Therefore, this model is taken as 'best' model, which we also use to examine the presence of non-linearities and interactions and to decompose a variety of different bond spread changes.

3.3. Examination of the presence of non-linearities and interactions

In contrast to linear regression models, machine learning methods in general are able to automatically learn non-linearities and interactions from the data. As all machine learning models outperform the linear regression model, we assume that their major performance can be attributed to such non-linearities and/or interactions. To investigate the assumption of non-linearities, we create partial dependence plots for all explanatory variables of all credit spread change models. Due to space constraints and better readability, we only show selected partial dependence plots of the model with target $\Delta CS^{\text{USD Corporate 5-7}}$ in Fig. 5.

Table 5
R^2 (in-sample) respectively SSR ratio (out-of-sample).

	In-sample Performance							Out-of-sample Performance						
	OLS	LPR	SVR	RF	BART	BRANN	LSTM	OLS	LPR	SVR	RF	BART	BRANN	LSTM
$\Delta CS^{\text{USD Corporate 1-3}}$	0.207	0.173	0.366	0.784	0.654	0.214	0.544	1.000	0.201	0.216	0.193	0.410	0.648	0.235
$\Delta CS^{\text{USD Corporate 3-5}}$	0.138	0.114	0.373	0.761	0.542	0.102	0.467	1.000	0.270	0.298	0.283	0.374	0.720	0.279
$\Delta CS^{\text{USD Corporate 5-7}}$	0.060	0.047	0.283	0.739	0.330	0.042	0.367	1.000	0.302	0.373	0.395	0.432	0.469	0.422
$\Delta CS^{\text{USD Corporate 7-10}}$	0.158	0.117	0.341	0.779	0.484	0.161	0.562	1.000	0.385	0.461	0.530	0.613	0.986	0.357
$\Delta CS^{\text{USD Corporate 10+}}$	0.189	0.089	0.373	0.787	0.338	0.180	0.452	1.000	0.509	0.579	0.608	0.712	0.883	0.521
$\Delta CS^{\text{USD Covered 1-3}}$	0.231	0.149	0.431	0.782	0.582	0.283	0.487	1.000	0.471	0.537	0.625	0.761	1.392	0.573
$\Delta CS^{\text{USD Covered 3-5}}$	0.255	0.126	0.426	0.792	0.518	0.236	0.500	1.000	0.596	0.670	0.600	0.756	0.997	0.717
$\Delta CS^{\text{USD Covered 5-7}}$	0.036	0.053	0.273	0.753	0.276	0.023	0.307	1.000	0.681	0.754	0.771	0.836	0.778	0.756
$\Delta CS^{\text{USD Covered 7-10}}$	0.256	0.205	0.403	0.805	0.731	0.308	0.582	1.000	0.274	0.368	0.343	1.066	1.539	0.348
$\Delta CS^{\text{USD Covered 10+}}$	0.190	0.153	0.288	0.797	0.607	0.270	0.313	1.000	0.558	0.579	0.556	0.818	0.921	0.571
$\Delta CS^{\text{EURO Corporate 1-3}}$	0.306	0.244	0.362	0.794	0.727	0.365	0.434	1.000	0.429	0.432	0.342	0.414	0.956	0.395
$\Delta CS^{\text{EURO Corporate 3-5}}$	0.338	0.247	0.421	0.798	0.711	0.399	0.484	1.000	0.521	0.504	0.390	0.385	1.007	0.425
$\Delta CS^{\text{EURO Corporate 5-7}}$	0.342	0.262	0.435	0.803	0.777	0.371	0.505	1.000	0.467	0.454	0.372	0.391	1.137	0.486
$\Delta CS^{\text{EURO Corporate 7-10}}$	0.323	0.235	0.408	0.796	0.731	0.327	0.439	1.000	0.424	0.418	0.349	0.407	0.692	0.378
$\Delta CS^{\text{EURO Corporate 10+}}$	0.269	0.229	0.438	0.805	0.702	0.267	0.465	1.000	0.518	0.503	0.438	0.540	1.052	0.419
$\Delta CS^{\text{EURO Covered 1-3}}$	0.031	0.044	0.168	0.691	0.422	0.028	0.082	1.000	0.535	0.602	0.603	0.677	0.881	0.547
$\Delta CS^{\text{EURO Covered 3-5}}$	0.047	0.057	0.233	0.734	0.333	0.041	0.140	1.000	0.719	0.751	0.765	0.790	0.884	0.796
$\Delta CS^{\text{EURO Covered 5-7}}$	0.047	0.048	0.249	0.751	0.290	0.035	0.173	1.000	0.680	0.728	0.766	0.731	0.792	0.726
$\Delta CS^{\text{EURO Covered 7-10}}$	0.064	0.056	0.254	0.742	0.334	0.045	0.152	1.000	0.477	0.475	0.579	0.650	0.704	0.547
$\Delta CS^{\text{EURO Covered 10+}}$	0.059	0.059	0.242	0.763	0.239	0.047	0.123	1.000	0.599	0.712	0.741	0.808	0.779	0.714
Average	0.177	0.135	0.338	**0.773**	0.516	0.187	0.379	1.000	**0.481**	0.521	0.512	0.629	0.911	0.511

Table 6
RMSE.

	In-sample Performance							Out-of-sample Performance						
	OLS	LPR	SVR	RF	BART	BRANN	LSTM	OLS	LPR	SVR	RF	BART	BRANN	LSTM
$\Delta CS^{\text{USD Corporate 1-3}}$	0.193	0.198	0.173	0.101	0.128	0.193	0.147	0.309	0.138	0.143	0.136	0.198	0.249	0.150
$\Delta CS^{\text{USD Corporate 3-5}}$	0.175	0.178	0.149	0.092	0.128	0.179	0.138	0.254	0.132	0.138	0.135	0.155	0.215	0.134
$\Delta CS^{\text{USD Corporate 5-7}}$	0.154	0.155	0.134	0.081	0.130	0.155	0.126	0.183	0.101	0.112	0.115	0.121	0.126	0.119
$\Delta CS^{\text{USD Corporate 7-10}}$	0.144	0.148	0.128	0.074	0.113	0.144	0.104	0.162	0.100	0.110	0.118	0.127	0.161	0.097
$\Delta CS^{\text{USD Corporate 10+}}$	0.129	0.136	0.113	0.066	0.116	0.129	0.106	0.143	0.102	0.109	0.111	0.121	0.134	0.103
$\Delta CS^{\text{USD Covered 1-3}}$	0.175	0.184	0.150	0.093	0.129	0.169	0.143	0.212	0.146	0.155	0.168	0.185	0.250	0.161
$\Delta CS^{\text{USD Covered 3-5}}$	0.208	0.226	0.183	0.110	0.167	0.211	0.171	0.209	0.161	0.171	0.162	0.181	0.208	0.177
$\Delta CS^{\text{USD Covered 5-7}}$	0.186	0.185	0.162	0.094	0.162	0.188	0.158	0.159	0.131	0.138	0.140	0.145	0.140	0.138
$\Delta CS^{\text{USD Covered 7-10}}$	0.208	0.215	0.186	0.106	0.125	0.201	0.156	0.199	0.104	0.121	0.117	0.206	0.247	0.117
$\Delta CS^{\text{USD Covered 10+}}$	0.459	0.470	0.431	0.230	0.320	0.436	0.423	0.368	0.275	0.280	0.275	0.333	0.353	0.278
$\Delta CS^{\text{EURO Corporate 1-3}}$	0.207	0.216	0.198	0.113	0.130	0.198	0.187	0.302	0.198	0.198	0.176	0.194	0.295	0.190
$\Delta CS^{\text{EURO Corporate 3-5}}$	0.182	0.194	0.171	0.101	0.121	0.174	0.161	0.276	0.199	0.196	0.173	0.171	0.277	0.180
$\Delta CS^{\text{EURO Corporate 5-7}}$	0.182	0.193	0.169	0.100	0.106	0.178	0.158	0.273	0.187	0.184	0.166	0.171	0.291	0.190
$\Delta CS^{\text{EURO Corporate 7-10}}$	0.173	0.184	0.162	0.095	0.109	0.172	0.157	0.267	0.174	0.172	0.157	0.170	0.222	0.164
$\Delta CS^{\text{EURO Corporate 10+}}$	0.115	0.118	0.101	0.059	0.073	0.115	0.098	0.214	0.154	0.152	0.142	0.157	0.219	0.138
$\Delta CS^{\text{EURO Covered 1-3}}$	0.128	0.127	0.119	0.072	0.099	0.128	0.125	0.088	0.065	0.069	0.069	0.073	0.083	0.065
$\Delta CS^{\text{EURO Covered 3-5}}$	0.105	0.105	0.094	0.056	0.088	0.106	0.100	0.087	0.074	0.075	0.076	0.077	0.082	0.078
$\Delta CS^{\text{EURO Covered 5-7}}$	0.108	0.108	0.096	0.055	0.093	0.108	0.100	0.089	0.073	0.076	0.078	0.076	0.079	0.076
$\Delta CS^{\text{EURO Covered 7-10}}$	0.095	0.095	0.084	0.050	0.080	0.096	0.090	0.093	0.065	0.064	0.071	0.075	0.078	0.069
$\Delta CS^{\text{EURO Covered 10+}}$	0.097	0.097	0.087	0.049	0.087	0.098	0.094	0.086	0.067	0.073	0.074	0.077	0.076	0.073
Average	0.171	0.176	0.154	**0.090**	0.125	0.169	0.147	0.199	**0.132**	0.137	0.133	0.151	0.189	0.135

The PDPs for the remaining explanatory variables can be found in the supplementary materials. Each of these figures depicts the possible values of the respective explanatory variable on the x-axis and the corresponding partial dependence values — i.e. the average prediction when all data points assume this value of the explanatory variable — on the y-axis. At the bottom of each PDP, a rug — that is to say an indicator for data points — is displayed on the x-axis to represent the variable distribution. The area between the 25th and the 75th quantile of the respective explanatory variable is highlighted in yellow, the supplementing area between the 5th and 95th quantile in blue.

Taking a look at the PDPs immediately shows that the functional relationship between the explanatory variables and the target is non-linear. Moreover, comparing the sign of the relationships between explanatory variables and target with those expected in the literature such as Collin-Dufresne et al. (2001) shows, that they mostly agree on the yellow area where the majority of data points is located. The same analysis is performed for all other types of credit spread changes including US corporate, US covered, EURO corporate and EURO covered bond spread changes of all considered maturities. We observe non-linearities in all of these cases though the functional form of the relationship between the explanatory variables and the credit spread changes varies slightly.

As mentioned before, besides using these PDPs to examine the functional form of the relationship between bond spread changes and explanatory variables, they can also be used to perform a hedging. We exemplary show the results of the factor hedging for selected explanatory variables and target $\Delta CS^{\text{USD Corporate 5-7}}$ in Fig. 6.

Taking a look at Fig. 6 shows that the factor hedgings $\alpha_i \cdot PDP^{CDS}_{x_i}$ — indicated by the red lines — are quite similar to the original credit spread partial dependence plots $PDP^{CS}_{x_i}$ — colored in black — within the yellow and blue plot areas. This indicates that a hedging using solely the explanatory variable under consideration is successfully possible.

Next, to examine the assumption of the presence of interactions, we investigate the H-statistic for all explanatory variables of all credit spread change models. Due to space constraints, we only show the H-statistic of the US corporate bond spread change model of maturity 5–7 years in Fig. 7. The plot depicts the H-statistics for all explanatory variables used in the model for the US corporate bond spread changes

Table 7
Direction-of-change accuracy.

	In-sample Performance							Out-of-sample Performance						
	OLS	LPR	SVR	RF	BART	BRANN	LSTM	OLS	LPR	SVR	RF	BART	BRANN	LSTM
$\Delta CS^{\text{USD Corporate 1-3}}$	0.512	0.533	0.704	0.883	0.696	0.525	0.650	0.556	0.500	0.556	0.569	0.611	0.556	0.542
$\Delta CS^{\text{USD Corporate 3-5}}$	0.496	0.608	0.713	0.871	0.700	0.508	0.683	0.417	0.542	0.556	0.528	0.431	0.417	0.458
$\Delta CS^{\text{USD Corporate 5-7}}$	0.579	0.596	0.700	0.929	0.683	0.558	0.762	0.389	0.486	0.431	0.472	0.444	0.444	0.472
$\Delta CS^{\text{USD Corporate 7-10}}$	0.567	0.592	0.725	0.883	0.675	0.575	0.733	0.514	0.403	0.444	0.486	0.500	0.500	0.500
$\Delta CS^{\text{USD Corporate 10+}}$	0.617	0.617	0.746	0.908	0.650	0.596	0.746	0.514	0.500	0.514	0.500	0.458	0.542	0.500
$\Delta CS^{\text{USD Covered 1-3}}$	0.617	0.554	0.683	0.921	0.717	0.608	0.688	0.458	0.528	0.514	0.444	0.528	0.458	0.528
$\Delta CS^{\text{USD Covered 3-5}}$	0.608	0.625	0.754	0.942	0.758	0.633	0.754	0.444	0.639	0.472	0.556	0.472	0.458	0.500
$\Delta CS^{\text{USD Covered 5-7}}$	0.529	0.604	0.667	0.883	0.650	0.483	0.683	0.403	0.458	0.458	0.569	0.486	0.375	0.458
$\Delta CS^{\text{USD Covered 7-10}}$	0.571	0.546	0.667	0.896	0.729	0.588	0.733	0.472	0.486	0.444	0.514	0.542	0.514	0.514
$\Delta CS^{\text{USD Covered 10+}}$	0.613	0.629	0.733	0.921	0.758	0.633	0.729	0.556	0.611	0.597	0.708	0.625	0.569	0.639
$\Delta CS^{\text{EURO Corporate 1-3}}$	0.637	0.642	0.762	0.883	0.717	0.658	0.708	0.625	0.583	0.708	0.597	0.639	0.611	0.569
$\Delta CS^{\text{EURO Corporate 3-5}}$	0.675	0.638	0.796	0.871	0.679	0.654	0.713	0.681	0.653	0.681	0.653	0.667	0.667	0.667
$\Delta CS^{\text{EURO Corporate 5-7}}$	0.637	0.654	0.779	0.863	0.721	0.646	0.729	0.667	0.681	0.653	0.694	0.722	0.681	0.667
$\Delta CS^{\text{EURO Corporate 7-10}}$	0.679	0.633	0.762	0.854	0.696	0.671	0.708	0.681	0.653	0.694	0.694	0.722	0.681	0.681
$\Delta CS^{\text{EURO Corporate 10+}}$	0.625	0.638	0.771	0.896	0.729	0.617	0.700	0.667	0.639	0.667	0.694	0.653	0.667	0.722
$\Delta CS^{\text{EURO Covered 1-3}}$	0.529	0.525	0.754	0.833	0.662	0.500	0.650	0.444	0.472	0.500	0.444	0.472	0.431	0.625
$\Delta CS^{\text{EURO Covered 3-5}}$	0.558	0.596	0.767	0.871	0.662	0.579	0.654	0.556	0.486	0.542	0.583	0.514	0.569	0.528
$\Delta CS^{\text{EURO Covered 5-7}}$	0.562	0.604	0.762	0.887	0.679	0.554	0.667	0.500	0.514	0.500	0.569	0.556	0.528	0.528
$\Delta CS^{\text{EURO Covered 7-10}}$	0.546	0.567	0.762	0.875	0.637	0.554	0.717	0.514	0.486	0.611	0.514	0.417	0.625	0.528
$\Delta CS^{\text{EURO Covered 10+}}$	0.567	0.529	0.767	0.871	0.604	0.554	0.679	0.500	0.486	0.583	0.528	0.486	0.500	0.472
Average	0.586	0.596	0.739	**0.887**	0.690	0.585	0.704	0.528	0.540	0.556	**0.566**	0.547	0.540	0.555

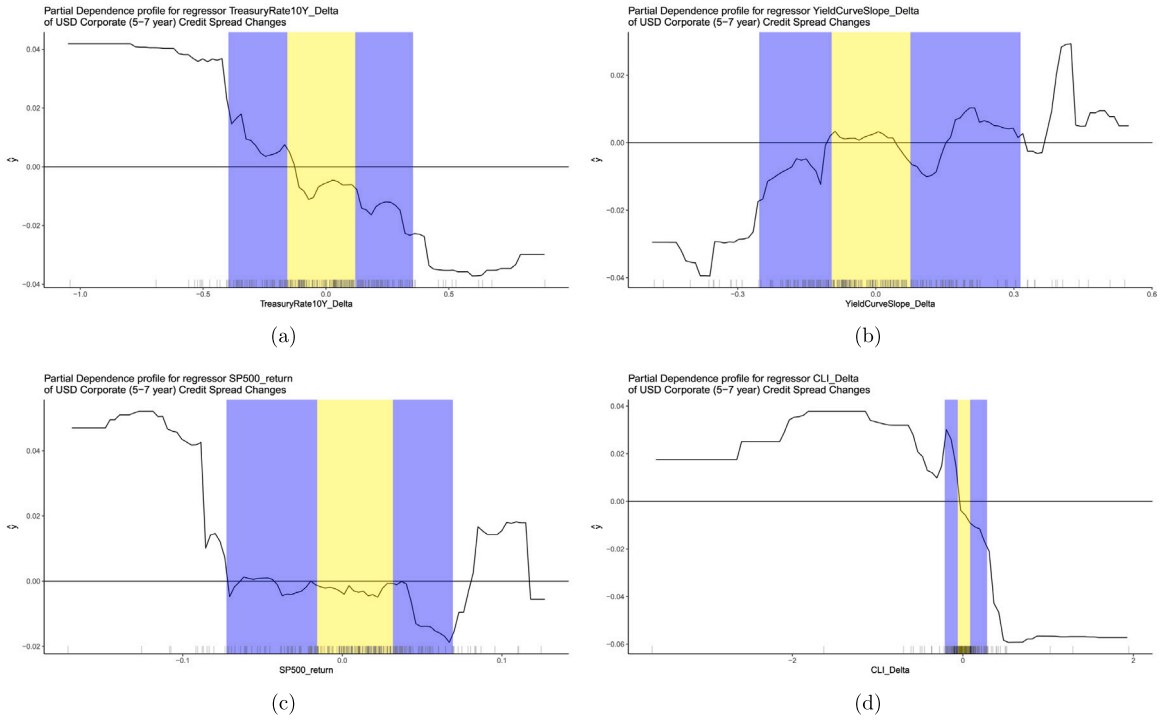

Fig. 5. Partial dependence plot for explanatory variables (a) Δr^{10}, (b) ΔSlope, (c) $\Delta \text{ret}^{SP500}$ and (d) ΔCLI. The yellow area indicates values of explanatory variables between 25% and 75% quantiles. The supplementing blue area indicates explanatory variables' values between 5% and 95% quantiles.

of maturity 5–7 years. As can be seen, the change in bid–ask spread exhibits the highest H-statistic, followed by the change in SKEW and the change in OECD's CLI. Overall, even the highest H-statistic scores are only around 0.3, implying that at most 30% of the total variation explained by the respective variable can be attributed to interactions with all other variables. Hence, all in all the evidence for interactions in US corporate bond credit spread change models of 5–7 year maturity is rather low.

This observation also holds for all other bond spread changes including US corporate, US covered and Euro Area corporate and covered bond spread changes of all maturities. Therefore, we conclude that the major reason for the out-performance of machine learning models compared to linear regression models can be attributed to non-linearities instead of interactions.

3.4. Decomposition of US corporate bond spread changes

After having fit the random forest model to our data, we apply the SHAP framework introduced in Section 2 to receive SHAP values for every observation (i.e. every month) and every considered explanatory variable. In Table 8, the resulting variable attributions of the model with target $\Delta CS^{\text{USD Corporate 5-7}}$ are exemplary visualized for the first observation (January 2002) of our data set.

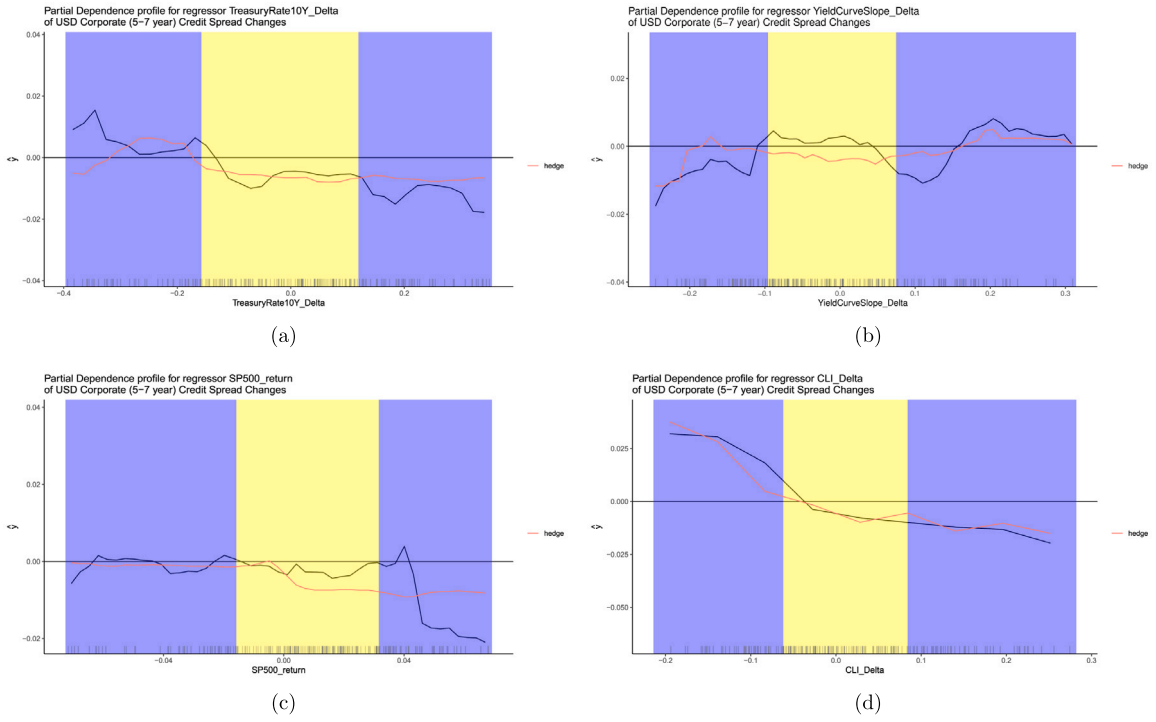

Fig. 6. Factor Hedging for explanatory variables (a) Δr^{10}, (b) ΔSlope, (c) $\Delta \text{ret}^{SP500}$ and (d) ΔCLI. The yellow area indicates values of explanatory variables between 25% and 75% quantiles. The supplementing blue area indicates explanatory variables' values between 5% and 95% quantiles.

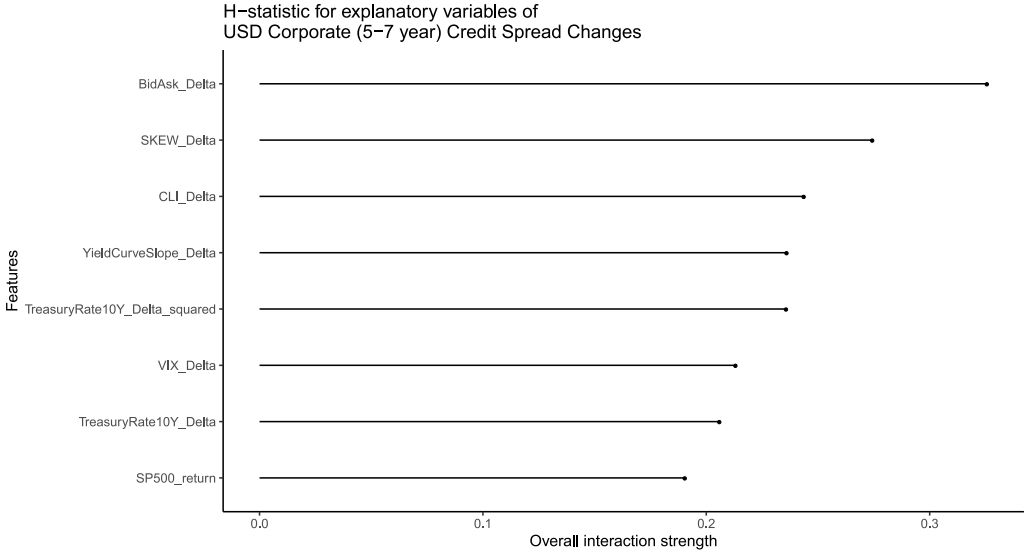

Fig. 7. H-statistic for all explanatory variables of USD corporate bond spread changes of 5–7 year maturity.

Table 8
Variable attributions.

Explanatory variable	Attribution
Δr^{10}	−0.0028
$(\Delta r^{10})^2$	−0.0029
ΔSlope	0.0031
ΔVIX	−0.0077
ret^{SP500}	−0.0032
ΔSKEW	−0.0091
ΔBidAsk	0.0009
ΔCLI	−0.0122

Plotting the values for the first observation listed in the above table results in the stacked bar for January 2002 in Fig. 8. To be precise, this bar is a stack of the eight explanatory variables' attributions to the spread change. This is now done for every observation in the data set. Visualizing the results yields the barplot shown in Fig. 8.

As can be seen in Fig. 8, there are extreme peaks and troughs in the USD corporate credit spread changes of 5–7 years maturity which seem to correspond to major crisis periods such as the financial crisis around 2008 and the Covid-19 pandemic in 2020. Note that this finding is in line with Guha and Hiris (2002), who empirically investigate the relationship between corporate bond spreads and the business cycle. According to their observations, the spread behaves counter-cyclically since it narrows during business cycle expansions and widens during contractions.

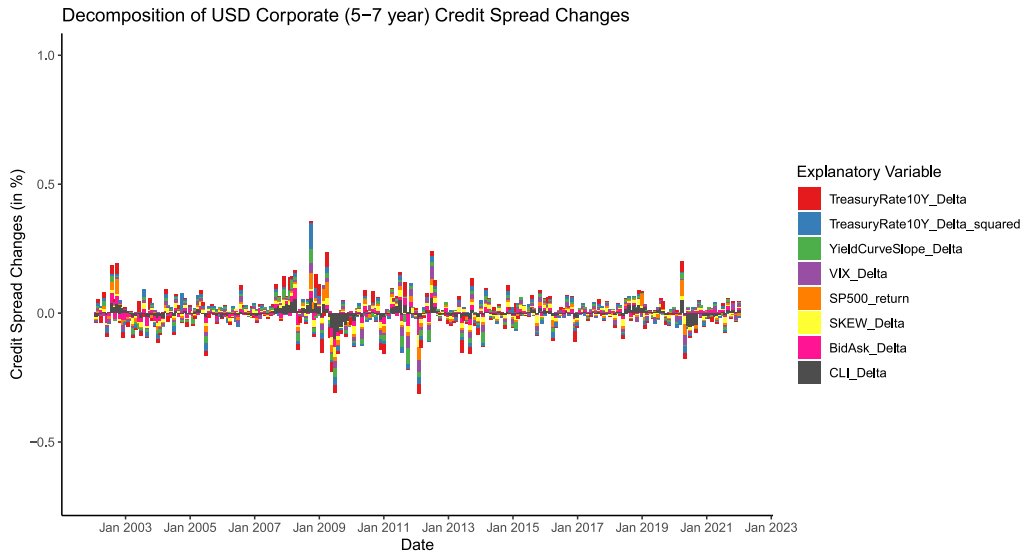

Fig. 8. Decomposition of USD corporate bond credit spread changes for 5–7 year maturity over the period from January 2002 to December 2021.

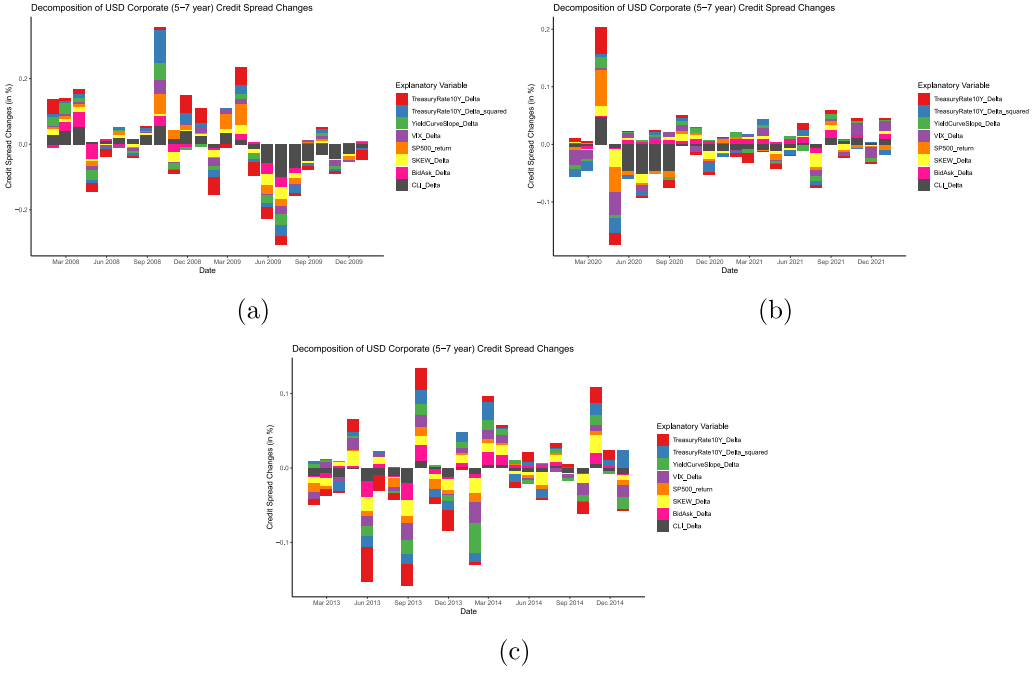

Fig. 9. Decomposition of USD corporate bond (5–7 year) credit spread changes over the period from (a) January 2008 to December 2009, (b) January 2020 to December 2021 and (c) January 2013 to December 2014.

Moreover, we find that the same holds true for USD corporate bond credit spread changes of all other maturities though the peaks and troughs are more extreme for shorter maturities and tend to get less intense when maturity increases. Hence, the predicted spread changes indicate that crisis have less impact on changes in long maturity bond spreads. We refer the interested reader to the supplementary materials for the visualization of the decomposition of the USD corporate bond credit spread changes of all other maturities.

As Fig. 8 is quite tiny, in Fig. 9 we zoom into the decomposition of the USD corporate (5–7 year) credit spread changes at three different periods. First, we take a look at the period of the financial crisis between January 2008 and December 2009. This crisis period is compared to the period between January 2020 and December 2021 capturing the Covid-19 pandemic. In addition, we compare the two crisis periods to the non-crisis period from January 2013 to December 2014.

The plot visualized in part (a) of Fig. 9 provides the decomposition of the modeled US corporate bond credit spread changes over the period between January 2008 and December 2009. It shows a strong peak in September 2008 implying that the US corporate bond credit spread strongly increased between the end of August and the end of September 2008. This sharp increase coincides with the bankruptcy of the investment bank Lehman Brothers and the resulting stock market crash in September 2008. This strong increase in the spread can be mostly attributed to squared changes in the 10-year Treasury rate, S&P 500 returns, changes in the yield curve slope and changes in CLI, which is in line with the fact that the CLI faced an extreme trough at that point in time. The next-largest peak can be observed in March 2009 when the market plummeted even more, leading to panicking investors. This strong increase in spread is mostly driven by changes in the 10-year Treasury rate and S&P 500 returns, which is in line with the fact

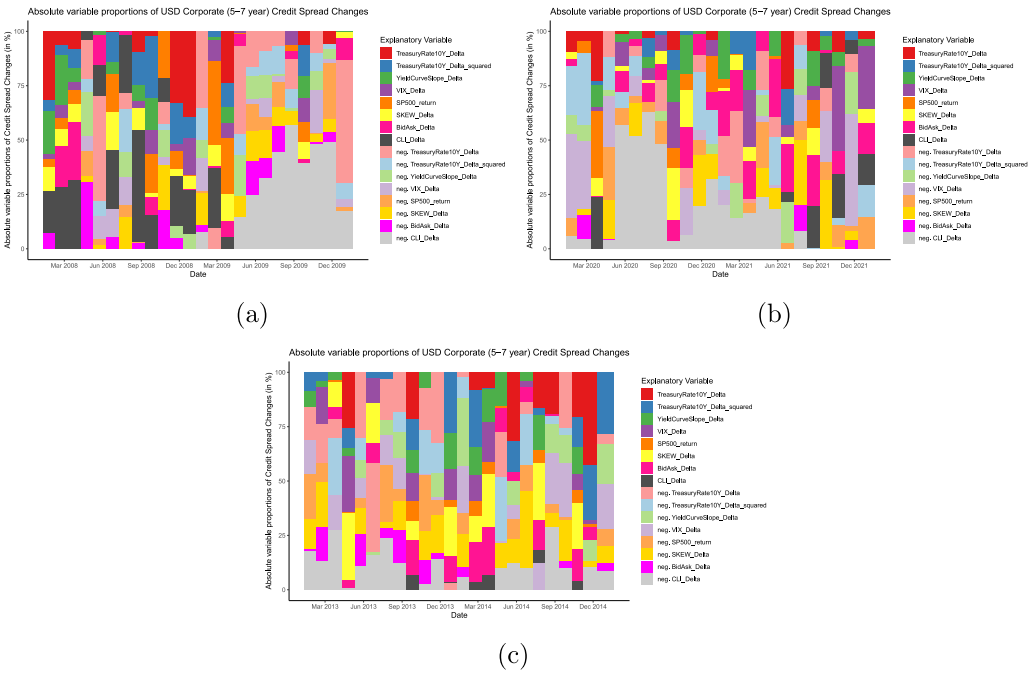

Fig. 10. Absolute variable proportions of the decomposition of USD corporate bond (5–7 year) spread changes over the period from (a) January 2008 to December 2009, (b) January 2020 to December 2021 and (c) January 2013 to December 2014.

that the S&P 500 faced an extreme trough at that time. In April 2009, the European Central Bank provided a swap line of $80 billion to the Federal Reserve, the British Central Bank granted them £60 billion, the Swiss central bank provided CHF 40 billion and the Japanese central bank 10 billion Yen. At that time, the spread starts to decline and does so until the end of the time period depicted in part (a) of Fig. 9. This change can be mostly attributed to changes in the CLI, which indeed constantly increases over this time period.

Comparison of these observations with the decomposition over the period of the Covid-19 pandemic depicted in part (b) of Fig. 9 yields quite similar results. There is an extreme peak in March 2020 corresponding to the rapid spread of the Covid-19 pandemic. Again, this strong increase in spreads can be attributed to changes in CLI and S&P 500 returns as well as to changes in the 10-year Treasury rate, which is in line with the strong drop in CLI. However, unlike during the financial crisis, financial markets recovered by far faster which is reflected by the fact that the spreads already start to constantly decrease the month after the strong peak. Again, the decline in spreads can be mostly attributed to changes in CLI which corresponds to the fact that CLI constantly recovers over this time period.

Comparison of both crisis periods with the non-crisis period between January 2013 and December 2014 visualized in part (c) of Fig. 9 yields that the magnitude of the spread changes is smaller during non-crisis times. Moreover, the changes in spreads are less driven by changes in CLI and more by changes in SKEW though the spread changes are more constantly attributed to the different explanatory variables.

To facilitate the recognition of the percentage amount each explanatory variable contributes to the spread change, we scale the attributions visualized in Fig. 9 to being between 0 and 100. More formally, let $b_{t1}, b_{t2}, b_{t3}, b_{t4}, b_{t5}, b_{t6}, b_{t7}$ and b_{t8} denote the attributions of the eight explanatory variables listed in Table 3 to the spread change ΔCS_t at time t respectively. Then, the absolute variable proportion $\hat{b}_{tj}$ of explanatory variable $j \in \{1,2,3,4,5,6,7,8\}$ at time t is given by

$$\hat{b}_{tj} = \frac{|b_{tj}|}{\sum_{k=1}^{8} |b_{tk}|}. \tag{7}$$

These absolute variable proportions are calculated for every observation in the data set. The resulting absolute proportions are visualized in Fig. 10. Note that variable shares, which contribute negatively to the spread changes are highlighted in the light colored version of their positive counterpart.

As can be seen in Fig. 10, this type of visualization strongly supports our findings that the positive peaks in the onset of the financial crisis and the Covid-19 pandemic can be mostly attributed to changes in CLI. Likewise, the constant decline in spreads in the offset of both crisis can again be mostly attributed to changes in CLI. Moreover, part (c) of Fig. 10 supports our claim that during non-crisis periods, the changes in spreads are way more evenly distributed across all explanatory variables.

3.5. Euro area versus US corporate bond spread change decomposition

In this subsection, we address the question regarding similarities and differences between US and Euro Area corporate bond spread changes by comparing their respective decompositions. To do so, the decomposition of the Euro Area corporate bond credit spread changes for 5–7 year maturity is visualized in Fig. 11. The visualization of the decomposition of the EURO corporate bond spread changes of all other maturities can be found in the supplementary materials. Note that these figures are obtained analogously to the US case described before.

As can be seen when comparing the decomposition of the Euro Area corporate bond credit spread changes with the US counterpart visualized in Fig. 8, the overall margin of the spread changes is similar. Likewise, also the overall evolution of the spread changes appears to be approximately the same. Particularly, we again observe extreme peaks during the financial crisis and the Covid-19 pandemic. Though, the margin of the peaks is larger than for US spread changes, especially in case of the Covid-19 peak. Moreover, we observe an additional peak in the middle of 2011 corresponding to the Euro-crisis. Overall, the margin of the aforementioned peaks again decreases with maturity though the decline is not as strong as in the US case.

To facilitate a detailed comparison, we only contrast the decomposition of EURO corporate bond (5–7 year) credit spread changes for the two crisis periods 2008–2009 and 2020–2021 and the non-crisis period between 2013 and 2014 visualized in Fig. 12 with the corresponding

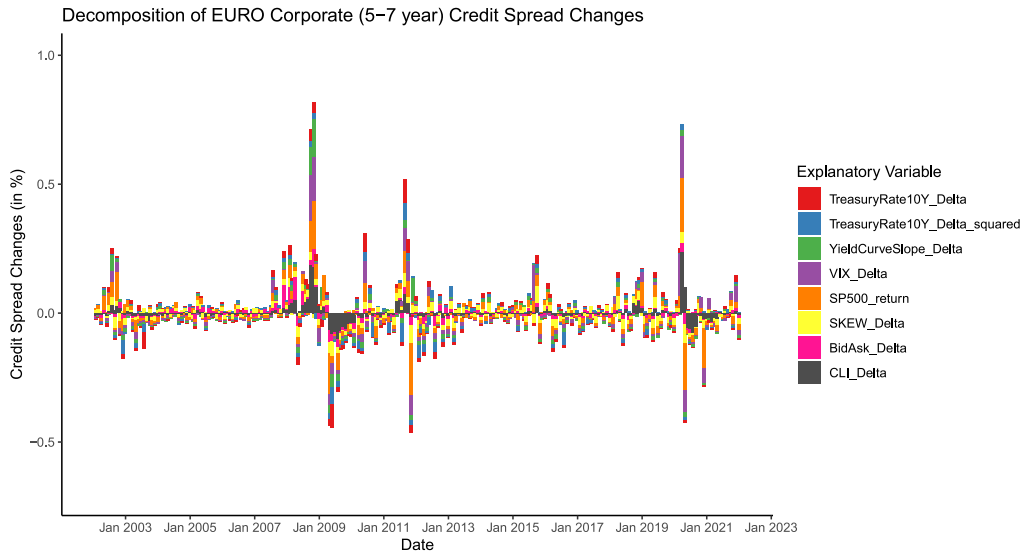

Fig. 11. Decomposition of EURO corporate bond credit spread changes for 5–7 year maturity over the period from January 2002 to December 2021.

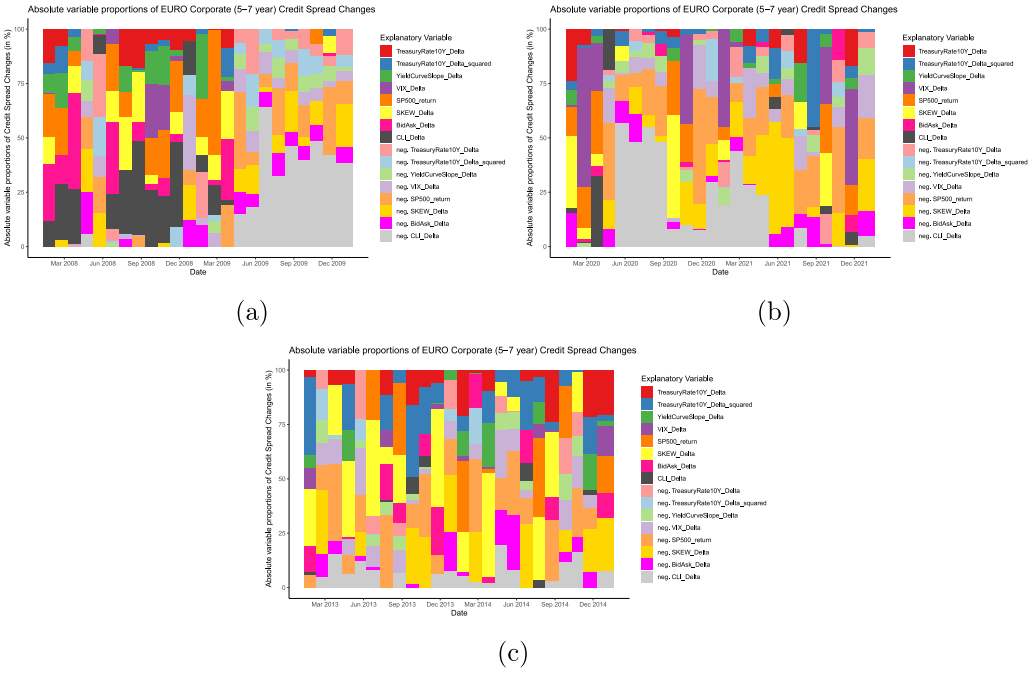

Fig. 12. Absolute variable proportions of the decomposition of EURO corporate bond (5–7 year) spread changes over the period from (a) January 2008 to December 2009, (b) January 2020 to December 2021 and (c) January 2013 to December 2014.

decompositions for the US corporate bond spread changes of maturity 5–7 years visualized in Fig. 10.

Comparing the absolute variable proportions of the decomposition of Euro Area corporate bond (5–7 year) spread changes depicted in Fig. 12 over the two crisis periods 2008–2009 and 2020–2021 and the non-crisis period 2013–2014 with their US counterparts visualized in Fig. 10, the overall evolution is quite similar. To be precise, in both cases the strong positive peaks in the onset of the financial crisis and the Covid-19 pandemic as well as the constant decline in spreads in the offset of both crisis can be mostly attributed to changes in CLI and S&P 500 returns. Though, there seems to be a stronger impact of changes in the VIX for Euro Area corporates during crisis periods compared to US corporates. Moreover, in contrast to the US case, part (c) of Fig. 12 shows that during non-crisis periods, the changes in Euro Area spreads are strongly driven by S&P 500 returns and changes in SKEW. In addition, all three plots in this figure convey the impression that Euro Area corporate bond spread changes are less driven by changes in 10-year Treasury rates than their US counterparts.

3.6. Corporate versus covered bond spread change decomposition

We apply our methodology to US and Euro Area covered bond spread changes and compare the resulting decompositions with their US and Euro Area corporate counterparts depicted in Figs. 8 and 11, respectively. For this purpose, the decompositions of US and Euro Area covered bond spread changes for maturity 5–7 years are visualized in Figs. 13 and 14, respectively. We refer the interested reader to the supplementary materials for visualizations of all other maturities.

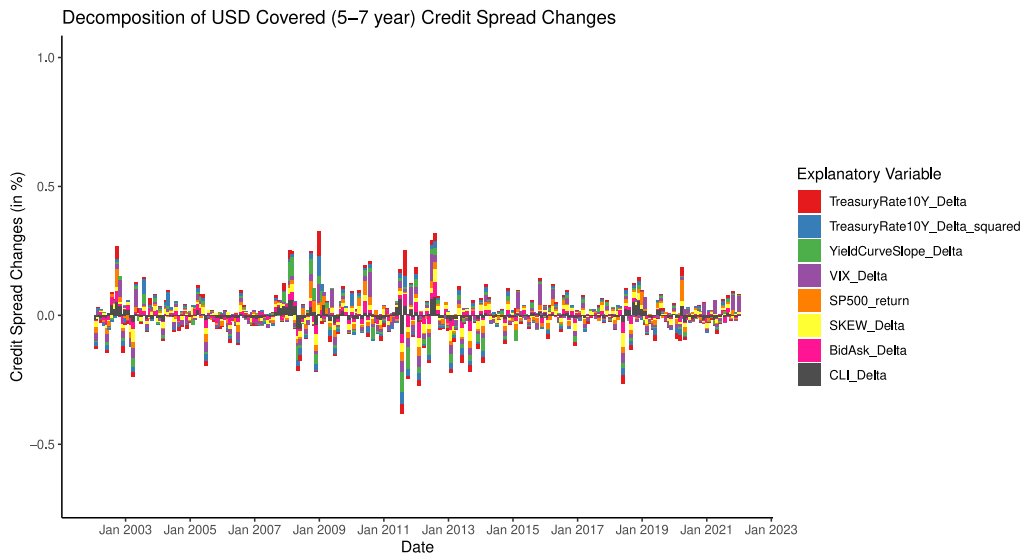

Fig. 13. Decomposition of USD covered bond credit spread changes for 5–7 year maturity over the period from January 2002 to December 2021.

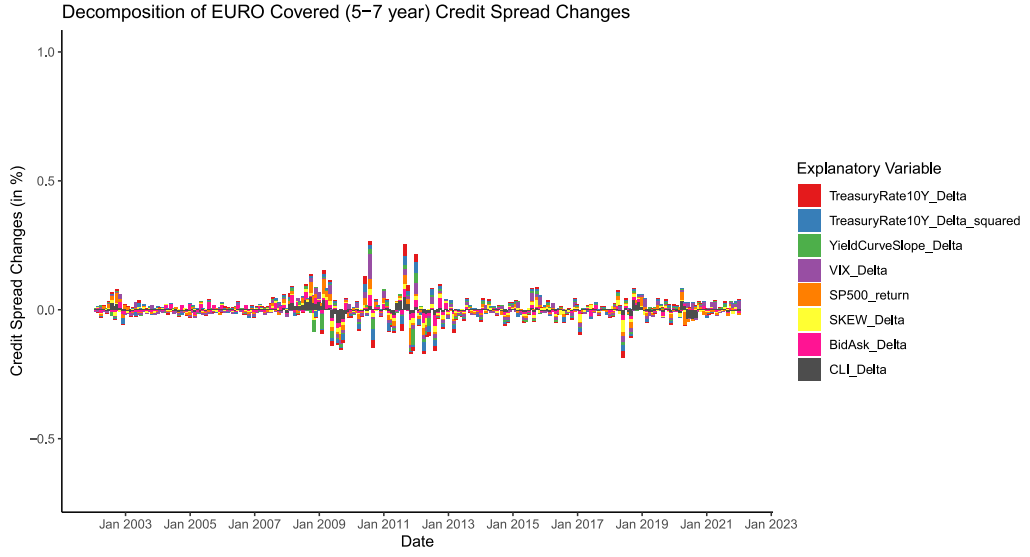

Fig. 14. Decomposition of EURO covered bond credit spread changes for 5–7 year maturity over the period from January 2002 to December 2021.

The decompositions of US and Euro Area covered bond spread changes are obtained analogously to the corporate bond spread change case. Further, the USD covered bond yield data is only available from January 2013. In addition, for maturities 5–7, 7–10 and 10+ year, the covered bond yield time series already end in December 2017, December 2015 and February 2016, respectively. For this reason, the missing data is extrapolated back to January 2002 and forth till December 2021 using linear regression. To do so, daily US corporate and sovereign and Euro Area covered and sovereign bond yields are used as possible explanatory variables and the best linear model is selected using bidirectional stepwise selection. The resulting model fit lies between 94% and 99% according to the coefficient of determination R^2.

Fig. 13 shows that the magnitude of the visualized US covered bond spread changes is in general smaller and less sensitive to crisis compared to the US corporate case. Moreover, the impact of crisis again seems to decrease with maturity. The magnitude of the Euro Area covered bond spread changes in Fig. 14 is even smaller and in general smallest across all spread changes that we compared during the course of our analysis. Moreover, the Euro Area covered bond spread changes seem to be only slightly affected by the major crisis. Interestingly, the impact of the Euro-crisis on the Euro Area covered bond spread changes even seems to be stronger than the impact of the financial crisis or the Covid-19 pandemic.

We again visualize the absolute variable proportions of US and Euro Area covered bond (5–7 year) credit spread changes each for the two crisis periods 2008–2009 and 2020–2021 and the non-crisis period between 2013 and 2014 in Figs. 15 and 16, respectively. These plots are compared to their corporate bond spread change counterparts visualized in Figs. 10 and 12, respectively.

Comparison between the USD covered bond spread change decompositions visualized in Fig. 15 and their USD corporate counterparts depicted in Fig. 10 yields that their overall evolution is quite similar. However, as can be seen in parts (a) and (b) of Fig. 15, the impact of the change in CLI during crisis is smaller compared to our previous analysis. In addition, especially during the time period between January 2020 and December 2021 that is visualized in part (b) of Fig. 15 there seems to be a partially enlarged impact of changes in VIX. Taking a look at part (c) of Fig. 15 yields that especially changes in SKEW have a larger impact in non-crisis periods compared to corporate bond spread changes.

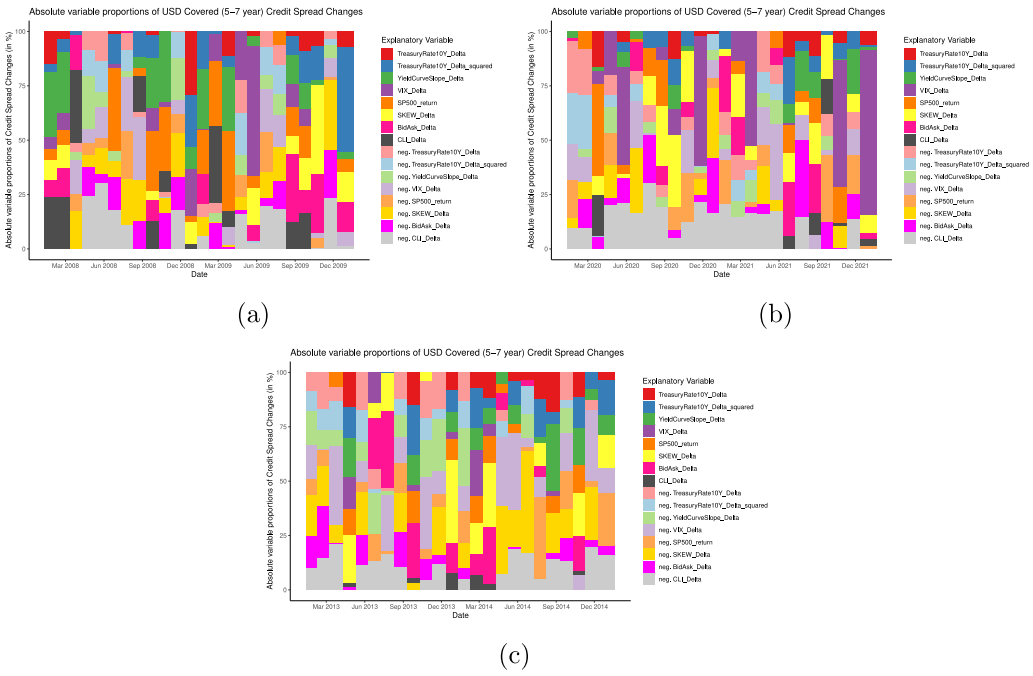

Fig. 15. Absolute variable proportions of the decomposition of USD covered bond (5–7 year) spread changes over the period from (a) January 2008 to December 2009, (b) January 2020 to December 2021 and (c) January 2013 to December 2014.

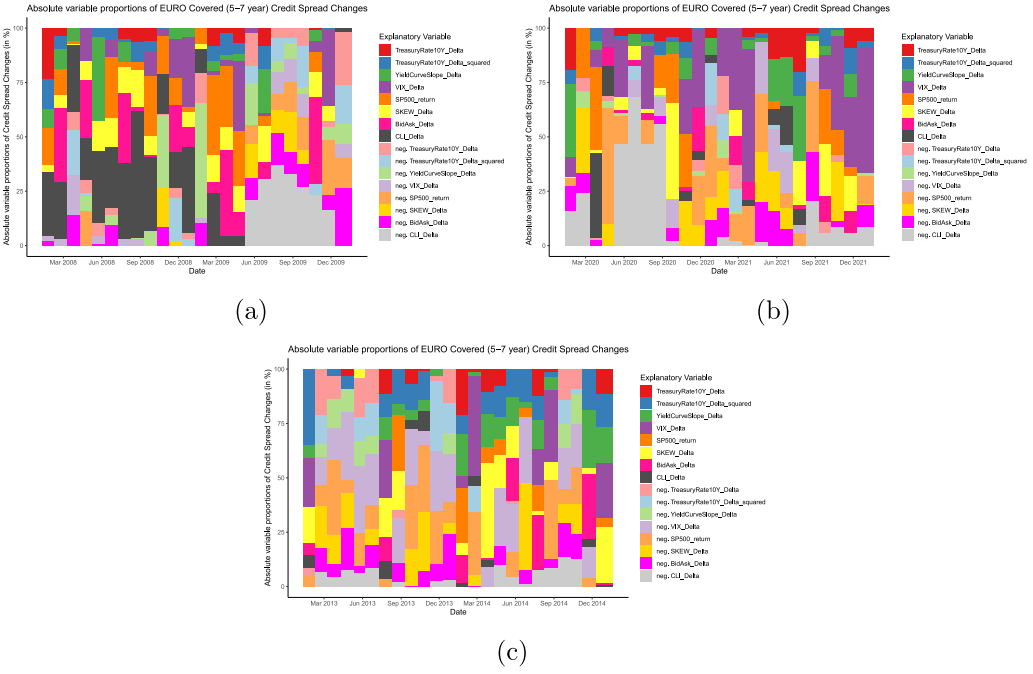

Fig. 16. Absolute variable proportions of the decomposition of EURO covered bond (5–7 year) spread changes over the period from (a) January 2008 to December 2009, (b) January 2020 to December 2021 and (c) January 2013 to December 2014.

In contrast to US covereds, the strong changes in Euro Area covered bond spreads during crisis can be mostly attributed to changes in CLI as shown in Fig. 16. In addition, similar to the US covered case, there seems to be an enlarged impact of changes in VIX during the time period between January 2020 and December 2021 that is visualized in part (b) of Fig. 16. Moreover, part (c) of Fig. 16 shows that changes in Euro Area covered bond spreads are quite evenly distributed across all explanatory variables during non-crisis periods, which is in line with the Euro Area corporate case.

3.7. Summary of credit spread change decomposition results

In this section, we tabularly summarize our findings regarding the decomposition of US and Euro Area corporate and covered bond credit spread changes in Table 9.

4. Conclusion and future research

In this paper, we aimed at modeling credit spread changes using a local polynomial regression model and different machine learning

Table 9
Summary of credit spread change decomposition results.

Similarities	Differences
US vs. Euro Area	
• similar time series evolution over time (strong peaks and troughs corresponding to major crisis periods such as financial crisis around 2008 and the Covid-19 pandemic in 2020) • impact of major crisis on both Euro Area and US spread changes decreases with maturity ⇒ the crisis have less impact on changes in long maturity bond spreads • spreads are more constantly attributed to the different explanatory variables during non-crisis periods • magnitude of spread changes is smaller during non-crisis periods	• margin of peaks during crisis periods is larger for Euro Area than for US spread changes (particularly for the Covid-19 crisis) • Euro Area spread changes show an additional peak in the middle of 2011 corresponding to the Euro-crisis • decline of the margin of the peaks with maturity is less strong in the Euro Area case compared to the US • in general differing impact of explanatory variables on spread changes • credit spread changes are mostly driven by OECD's CLI and S&P 500 returns in the US and Euro Area case but also by the 10-year Treasury rate in the US case and by the VIX in the Euro Area case during crisis periods • during non-crisis periods, US spread changes are mostly driven by changes in the SKEW whereas Euro Area spread changes are mostly driven by changes in the SKEW and S&P 500 returns
corporate vs. covered	
• similar time series evolution over time (strong peaks and troughs corresponding to major crisis periods such as financial crisis around 2008 and the Covid-19 pandemic in 2020) • impact of crisis decreases with maturity • magnitude of spread changes is smaller during non-crisis periods • spreads are more constantly attributed to the different explanatory variables during non-crisis periods	• corporate bond spread changes are in general wider and more sensitive to crisis compared to covered bond spread changes • in general differing impact of explanatory variables on spread changes • changes in VIX have larger impact on covereds than on corporate bond spread changes during crisis periods whereas in non-crisis periods changes in the SKEW are more important for covered bond spread changes than for their corporate counterparts

methods to analyze the presence of non-linearities and interactions as potential reasons for the failure of standard factors within linear regression models. For this purpose, we started with applying a selection of the most common machine learning models to a variety of different US and Euro Area credit spread changes to measure their performance. We find that all machine learning models on average outperform the linear regression in terms of all four considered performance measures both in-sample as well as out-of-sample. In addition, most machine learning approaches also outperform the local polynomial regression model. We identify the random forest as the model that fits the credit spread changes significantly best.

The random forest model is thus used for the examination of the presence of non-linearities and interactions as well as for the decomposition of the credit spread changes into the explanatory variables. To facilitate the analysis of non-linearities and interactions, we propose the usage of two popular explainable AI methods called partial dependence plots and H-statistic. We find that the major reason for the out-performance of machine learning models compared to linear regression can be attributed to the presence of non-linear relationships and less to interactions. Moreover, we provide evidence that PDPs can be successfully used to perform a factor hedging.

We propose the usage of SHAP values for the decomposition of credit spread changes. This method is applied to US and Euro Area corporate and covered bond spread changes of different maturities. We find that the overall movement of US corporate and covered bond spread changes mostly coincides with their Euro Area counterparts though the impact of the different explanatory variables on the changes might partly differ. In particular, we observe that onsets and offsets of crisis periods are well explained by changes in OECD's composite leading indicator and S&P 500 returns. Moreover, USD bond spread changes seem to be strongly driven by changes in the 10-year Treasury rate during crisis periods whereas in the Euro Area case, changes in the VIX are more important. During non-crisis periods, US spread changes are increasingly driven by changes in the SKEW whereas Euro Area spread changes are driven by changes in SKEW and S&P 500 returns. Likewise, corporate and covered bond spread changes overall evolve quite similar though the impact of the different explanatory variables on the changes differs slightly. The impact of the change in CLI during crisis is way smaller for US covered bond spread changes. Instead, changes in the VIX have a larger impact on covered bond spread changes during crisis periods. Moreover, in non-crisis periods, changes in SKEW seem to be more important for covered bond spread changes than for their corporate counterparts.

Besides the contents of this paper, there are some topics that can be further studied. First, we consider a selection of the most popular explanatory variables for modeling credit spread changes. This selection of features can be extended by considering further variables that have not yet been examined in the literature such as crisis indicators or probabilities (see e.g. Hauptmann et al. 2014). Second, besides our selection of models, there exist other promising ML models such as new architectures of neural network models or boosting approaches that can be tested. Third, we have proposed the usage of selected XAI techniques for examining the presence of non-linearities and interactions as well as for decomposing the credit spread changes. In recent years, the number of XAI methods that can be used for these purposes is ever growing. Hence it might be interesting to compare our results to those obtained by the application of further methods such as Variable Interaction Networks, Greenwell interactions, permutation-based feature importance, LIME, or breakDown (see e.g. Molnar 2022). Finally, an application of the proposed framework for the allocation of a bond portfolio is a next research topic.

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.

Data availability

The authors do not have permission to share data.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.irfa.2024.103315.

References

Amihud, Y., & Mendelson, H. (1986). Asset pricing and the bid-ask spread. *Journal of Financial Economics*, *17*, 223–249.
Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, *57*, 289–300.
Bianchi, D., Büchner, M., & Tamoni, A. (2021). Bond risk premiums with machine learning. *The Review of Financial Studies*, *34*, 1046–1089.
Breitenfellner, B., & Wagner, N. (2012). Explaining aggregate credit default swap spreads. *International Review of Financial Analysis*, *22*, 18–29.

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*, 1509–1531.

Chipman, H., George, E., & McCulloch, R. (2010). BART: Bayesian additive regression trees. *Annals of Applied Statistics*, *4*, 266–298.

Cleveland, W. S., Grosse, E., & Shyu, W. M. (1992). *Local regression models*. Wadsworth & Brooks/Cole, chapter 8.

Collin-Dufresne, P., Goldstein, R., & Martin, J. S. (2001). The determinants of credit spread changes. *The Journal of Finance*, *56*, 2177–2207.

Demsar, J. (2006). Statistical comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research*, *7*, 1–30.

Feldhütter, P., & Poulsen, T. K. (2018). What determines bid-ask spreads in over-the-counter markets? Available at SSRN: http://dx.doi.org/10.2139/ssrn.3286557.

Friedman, J. (2001). Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, *29*, 1189–1232.

Friedman, J., & Popescu, B. (2008). Predictive learning via rule ensembles. *Annals of Applied Statistics*, *2*, 916–954.

Friewald, N., & Nagler, F. (2019). Over-the-counter market frictions and yield spread changes. *The Journal of Finance*, *74*, 3217–3257.

Gers, F., Schmidhuber, J., & Cummins, F. (2000). Learning to forget: Continual prediction with LSTM. *Neural Computation*, *12*, 2451–2471.

Greene, W. (2003). *Econometric analysis*. Prentice Hall.

Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies*, *33*, 2223–2273.

Guha, D., & Hiris, L. (2002). Modeling credit spreads: An application to the sterling eurobond market. *International Review of Financial Analysis*, *11*, 219–227.

Hauptmann, J., Hoppenkamps, A., Min, A., Ramsauer, F., & Zagst, R. (2014). Forecasting market turbulence using regime-switching models. *Financial Markets and Portfolio Management*, *28*, 139–164.

He, Z., Khorrami, P., & Song, Z. (2022). Commonality in credit spread changes: Dealer inventory and intermediary distress. *The Review of Financial Studies*, hhac004.

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, *9*, 1735–1780.

Javadi, S., Nejadmalayeri, A., & Krehbiel, T. L. (2018). Do FOMC actions speak loudly? Evidence from corporate bond credit spreads. *Review of Finance*, *22*, 1877–1909.

Kaviani, M. S., Kryzanowski, L., Maleki, H., & Savor, P. (2020). Policy uncertainty and corporate credit spreads. *Journal of Financial Economics*, *138*, 838–865.

Kim, T. S., Park, J. W., & Park, Y. J. (2017). Macroeconomic conditions and credit default swap spread changes. *Journal of Futures Markets*, *37*, 766–802.

Lundberg, S., Erion, G., & Lee, S. I. (2018). Consistent individualized feature attribution for tree ensembles. CoRR abs/1802.03888.

Lundberg, S., & Lee, S. I. (2017). A unified appproach to interpreting model predictions. In *NIPS' 17, Preceedings of the 31st international conference on neural information processing systems* (pp. 4768–4777).

Manzoni, K. (2002). Modeling credit spreads: An application to the sterling eurobond market. *International Review of Financial Analysis*, *11*, 183–218.

Martell, R. (2008). Understanding common factors in domestic and international bond spreads. *Review of Finance*, *12*, 365–389.

Molnar, C. (2022). *Interpretable machine learning: A guide for making black box models explainable* (2nd ed.). christophm.github.io/interpretable-ml-book/.

Murphy, K. (2012). *Machine learning: A probabilistic perspective*. MIT Press.

Sariev, E., & Germano, G. (2019). Bayesian regularized artificial neural networks for the estimation of the probability of default. *Quantitative Finance*, *20*, 311–328.

Son, Y., Byun, H., & Lee, J. (2016). Nonparametric machine learning models for predicting the credit default swaps: An empirical study. *Expert Systems with Applications*, *58*, 210–220.

Xiong, R., Cai, H., Diego-Guerra, I., Lu, Y., Xu, X., & Yin, Y. (2019). Forecasting credit spreads: A machine learning approach. Semantic Scholar.

Zhou, Z. H. (2012). *Ensemble methods foundations and algorithms*. Chapman and Hall/CRC.