



HAL
open science

Dissecting the explanatory power of ESG features on equity returns by sector, capitalization, and year with interpretable machine learning

Jérémi Assael, Laurent Carlier, Damien Challet

► To cite this version:

Jérémi Assael, Laurent Carlier, Damien Challet. Dissecting the explanatory power of ESG features on equity returns by sector, capitalization, and year with interpretable machine learning. *Journal of Risk and Financial Management*, 2023, 16 (3), pp.159. 10.3390/jrfm16030159 . hal-03791538v2

HAL Id: hal-03791538

<https://hal.science/hal-03791538v2>

Submitted on 16 Dec 2022 (v2), last revised 7 Apr 2023 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Dissecting the explanatory power of ESG features on equity returns by sector, capitalization, and year with interpretable machine learning

Jérémi Assael^{1,2}, Laurent Carlier², and Damien Challet¹

¹Chair of Quantitative Finance, MICS Laboratory, CentraleSupélec,
Université Paris-Saclay, Gif-sur-Yvette, France

²BNP Paribas Corporate & Institutional Banking, Global Markets
Data & Artificial Intelligence Lab, Paris, France

Abstract

We systematically investigate the links between price returns and Environment, Social and Governance (ESG) features in the European market. We propose a cross-validation scheme with random company-wise validation to mitigate the relative initial lack of quantity and quality of ESG data, which allows us to use most of the latest and best data to both train and validate our models. Boosted trees successfully explain a part of annual price returns not accounted by the market factor. We check with benchmark features that ESG features do contain significantly more information than basic fundamental features alone. The most relevant sub-ESG feature encodes controversies. Finally, we find opposite effects of better ESG scores on the price returns of small and large capitalization companies: better ESG scores are generally associated with larger price returns for the latter, and reversely for the former.

Keywords: ESG features; ESG data; sustainable investing; interpretable machine learning; model selection; asset management; equity returns;

JEL Classification: C51; C52; C55; G11; G12; G41;

1 Introduction

Investing according to how well companies do with respect to their environmental, social and governance scores has become very appealing for a growing number of

investors. Beyond moral criteria, such kind of investments may increase the value of high-ESG scoring companies, which will attract even the non-ESG minded investor, thereby starting a virtuous circle both for the investors and for the beneficiaries of high ESG scores. It may also lead to successful impact investing whereby an investor generates positive environmental or societal impact while targeting a specific level of return (Townsend, 2020; Grim and Berkowitz, 2020).

From a quantitative point of view, ESG features raise the question of their information content. Friede et al. (2015) aggregated the results of more than 2200 studies: 90% of them showed a non-negative relation between ESG and corporate financial performance measures, a majority displaying a positive performance. However, more recently, Cornell and Damodaran (2020); Breedts et al. (2019); De Franco et al. (2020) reach less clear-cut conclusions. Our aim is to clarify this issue. The real question indeed is whether ESG features contain additional information with respect to well-known factors. This question is also becoming increasingly important with the emergence of new directives and regulations like the second Markets In Financial Instruments Directive (MiFID II), the Sustainable Finance Disclosure Regulation (SFDR) or the European Union taxonomy, concerning disclosure of information on sustainable development and sustainable investing.

The current situation is mostly caused by challenges specific to ESG data: i) they are quite sparse before 2015, as the interest of even computing such scores is quite recent; ii) they are usually updated yearly; iii) the way they are computed often changes as a function of time and may depend on the way companies disclose data; iv) human subjectivity may be involved to a large extent in the computation of the scores.

This work is devoted to overcoming these challenges by using high-quality data and by testing several model validation schemes, including random company-wise cross-validation that makes it possible to use most of the latest data to train models. The resulting methodology is robust enough to yield clear answers about the role of capitalization, sector, and time in the explanatory power of ESG factors.

2 Literature review and uncertainties

2.1 Asset selection, investment strategies, and portfolios

According to Chen and Mussalli (2020), ESG integration into investment strategies mainly consists in integrating the investors values into their own strategies. The scientific literature describes three main ways to achieve it: filtering companies based on their ESG scores, directly looking for alpha in ESG data, or measuring

ESG impact on other risk factors.

ESG scores, for example controversies, provide a systematic way to filtering "sin industries" related to tobacco, alcohol, pornography, weapons. . . Schofield et al. (2019) select companies with ESG scores above certain thresholds. While this method yields good portfolios ESG-wise, Alessandrini and Jondeau (2020) argue that this may lead to underperforming portfolios because this selection reduces the investment universe and because these "sin industries" can have higher returns because of their very exclusion.

Chen and Mussalli (2020) propose a Markowitz-like optimization method by defining an ESG-compatible efficient frontier. Similarly, Hilario-Caballero et al. (2020) add a third term to the mean-variance cost function, the portfolio exposure to carbon risk, and use a genetic algorithm to solve this three-criterion optimization problem. This method is equivalent to optimizing an ESG criteria under the constraint of specific risk and returns levels (Schofield et al., 2019), who also note that the resulting portfolio can have a good global ESG score while containing assets with bad ones.

Finally, Alessandrini and Jondeau (2020) elaborate on "smart beta" strategies, in which investors build portfolios whose assets are not weighted according to their market capitalization but rather to their exposure to some specific risk factors. Bacon and Ossen (2015) explain that integrating ESG into investment strategies can be simply achieved by tilting the asset weights according to their ESG scores, while controlling the portfolio exposure to other risk factors. This procedure raises the question whether ESG is a new risk factor or if optimizing ESG scores amounts to exposing the portfolio to well-known ones. It is indeed a crucial point to explore when attempting to improve portfolio performance with ESG scores (Anson et al., 2020): instead of trying to obtain a premium by finding a suitable ESG factor, it is more judicious to understand the impact of ESG data on the exposure to well-known risk factors.

2.2 ESG scores: risk and returns

Reaching a consensus on the nature of the links between ESG and returns is hard. Friede et al. (2015) aggregated more than 2000 studies on the topic: 41% did not find any ESG impact on returns, 48% found these impacts to be positive and 9% negative. Alessandrini and Jondeau (2020) and Anson et al. (2020) stress the fact that filtering a portfolio on ESG scores leads to an improved durability of the investment but does not yield a positive alpha. However, they did not find any proof of negative alpha either, thus there may be no added value in integrating ESG data into portfolio construction from an alpha point of view.

Plagge and Grim (2020) find no statistically significant under- or over-performance of different equity funds specialized in ESG investing and argue that since the ESG

scores are not of economic nature, they should not have any impact on the portfolios, and that any information contained in ESG data should already be contained in other risk factors.

This lack of consensus on the links between ESG and returns may be due to the use of different assessment methodologies (including ESG scores from different data providers), or to a wrong use of the ESG scores (Anson et al., 2020). Indeed, De Franco et al. (2020) show that because ESG data have a very low signal-to-noise ratio, using aggregated ESG scores leads to a high loss of information. It is then necessary to use more granular scores to obtain more meaningful results. Moreover, they emphasize that the links between ESG and returns are highly dependent on the considered business industry and region. Cappucci (2018) finds that ESG scores lack information on asset price returns and that a better indicator of returns is the progress made by companies in the different ESG sub-fields.

Only a few papers are devoted to the relationship between ESG scores and risk. Guo et al. (2020) train a deep learning model to predict a company's volatility using ESG news. Chen and Mussalli (2020) show that focusing on ESG investments can reduce the risk of under-performance as companies with good ESG scores can be less exposed to both systemic and idiosyncratic risks.

2.2.1 Risk factors

Many studies, such as Renshaw (2018), find that ESG scores and well-known risk factors are partially redundant. Anson et al. (2020) and Konqui et al. (2019) study the variation of portfolio exposure to well-known risk factors when one integrates ESG data in portfolio construction: the impact varies according to geographical regions, which reduces the significance of global studies. Similarly, Alessandrini and Jondeau (2020) explain that the discrepancies in ESG portfolio performances in different regions and industries can be attributed to different exposures to risk factors. Furthermore, Breedt et al. (2019) argue that most of the financial performance of a portfolio can be explained by well-known factors and that the residuals cannot be explained by any other factors. For Breedt et al. (2019), the environmental and social aspects of ESG are noise and the governance part is strongly correlated to the quality factor; however, enriching ESG data with other types of information, or preprocessing it, can bring added value. In the same vein, Bacon and Ossen (2015) decorrelate the ESG scores from the other risk factors before integrating them into strategies and are able to obtain added value from ESG scores.

2.2.2 Materiality of ESG data

For a better ESG integration, it is important to understand which ESG indicators are the most material, i.e., have the largest impact on the financial performance of a

company. According to Anson et al. (2020) and De Franco et al. (2020), materiality is highly dependent on the chosen asset class, region and industry. Bacon and Ossen (2015) build a materiality matrix using the LASSO method (Tibshirani, 1996). Their matrix is specific to an industry and shows the magnitude of impact of a specific ESG indicator on a company’s financial performance versus the probability of this indicator having an impact.

2.2.3 Temporality of ESG data

Alessandrini and Jondeau (2020) warn that their results were obtained in a period when a large amount of money was poured into ESG funds, which could have increased their respective performance. De Franco et al. (2020) also remind that their study was realised between 2009 and 2018 during a period when the market was particularly bullish, which may affect the overall strength of ESG-based funds. For Roncalli and Mortier (2019), the impact of ESG scores differs not only by region and by industry but also according to the strategy test time period. That is why Renshaw (2018) argues that any methodology which treats the historical ESG data in the same way for every period is likely not relevant. A solution is to use back-testing on several time periods, with several universes, to validate the results (Anson et al., 2020). Finally, De Franco et al. (2020) and Plagge and Grim (2020) apply the efficient markets theory in the context of ESG investing: it is possible that, investor awareness rises as a function of time, the information included in ESG data is included in the assets prices, leading to a loss of predictive power of ESG features and thus of the embedded alpha.

2.3 Our contribution

In this paper, we study the relationships between ESG and returns and find statistically significant results using machine learning methods on a geographically restricted universe of companies, with ESG data with the least amount of human intervention, and at the right level of ESG score granularity. We include the time evolution dimension in the models we train, propose a cross-validation scheme that allows us to retain the most recent (and relevant) data and analyse how each ESG score contributes to the overall performance of our algorithms. Finally, we propose a new way to build a materiality matrix based on the explainability of the chosen machine learning method.

3 Datasets

3.1 Financial data

We use the following data sets:

- Stock prices. We use daily close prices, adjusted for dividends and foreign exchange rates. BNP Paribas internal data sources.
- Market capitalization. BNP Paribas internal data sources.
- Fama-French market, size and value factors: these factors are taken from the online French data library (Fama and French, 2021). They are all computed according to the Fama and French methodology exposed in Fama and French (1993).
- Risk-free rate: this data is also taken from Fama and French (2021) and computed according to the Fama and French method.

In addition, metadata like the TRBC Sector at levels 1, 2 and 3 and the country of incorporation are used and come from Refinitiv data sources.

3.2 ESG Data

ESG data are provided by Refinitiv. Their database alleviates some of the challenges listed above:

1. The coverage of the dataset is sufficient to extract meaningful results. Figure 1 shows the number of samples in the geographical regions as defined by Fama and French (Fama and French, 2021): Europe, North America, Japan, Asia-Pacific excluding Japan and Emerging Countries. Refinitiv ESG data start in 2002 and the number of samples per year increased by several folds until 2019, as shown in Fig. 2. The drop in 2020 is due to the fact that not all the ESG scores had been computed by Refinitiv when we had access to the dataset (many companies had not yet published enough data). It will not be an important issue as training does not take place in 2020 but it may have an impact on the test results of this particular year.
2. Scores are built with a well documented methodology explained in Refinitiv (2021). Every ESG score ranges between 0 and 1, 1 being the best score. In addition, the same methodology is used throughout the years, yielding consistent data.
3. Human intervention is limited to some quality checks.




Figure 1: Number of samples in each Fama-French region in the Refinitiv ESG dataset.

4. Scores can be updated up to 5 years after first publication, which is beneficial in an explanatory setting, as the data become more accurate. In a purely predictive setting however, this adds noise and look-ahead bias as we do not have point-in-time data, i.e., we do not know the initial and intermediate ESG estimates.

Refinitiv ESG data includes samples from different regions of the world. Each region has specific regulatory frameworks and ESG transparency rules. This is why this paper focuses on the European region and includes all the companies in the Refinitiv ESG dataset whose country of incorporation is in Europe or in a European-dependent territory.

The European ESG dataset contains 20509 samples for 2429 companies uniquely identified by their ISIN. The time evolution of the number of samples per year is reported in Fig. 3. All the sectors have enough data, with the notable exception of the Academic and Educational Services sector (see Fig. 4).

4 Methods

4.1 Problem settings

Our goal is to understand how and what ESG features participate in the formation of price returns. More precisely, we examine whether ESG features bring additional information with respect to well-known factors. In a multi-factor model, one writes




Figure 2: Time evolution of the number of samples in the Refinitiv ESG dataset.




Figure 3: Time evolution of the number of samples per year in the Refinitiv ESG dataset - Europe.




Figure 4: Number of samples for TRBC L1 Sector in the Refinitiv ESG dataset - Europe.

at time t

$$r_{i,t} = r_{f,t} + \sum_k w_{i,k} F_{k,t} + \alpha_i + \epsilon_{i,t} \quad (1)$$

where $r_{i,t}$ is return of asset i , $r_{f,t}$ the risk-free rate, $F_{k,t}$ the value of factor k at time t and $w_{i,k}$ is the factor loading; the idiosyncratic parts are α_i , the unexplained average return, and the zero-average residuals $\epsilon_{i,t}$. In this work, we use the CAPM (Capital Asset Pricing Model) model and its extension, the Fama-French 3-factor model that includes market (r_m), size (SMB), and value (HML) factors (Fama and French, 1993).

When data is abundant and of constantly high quality, the determinants of the idiosyncratic part can be found with standard tools. However, in our case, we must settle for a less ambitious goal. We investigate indeed here the most basic question: can ESG features help explain the sign of the idiosyncratic part of price returns? Mathematically, one needs to explain

$$Y_{i,t} = \frac{1 + \text{sign}(\alpha_i + \epsilon_{i,t})}{2} \quad (2)$$

with the candidate features.

This work takes a machine learning approach to this problem and treats it as a classification problem: $Y_{i,t}$ defines two classes as it can take two values. Thus, for each possible couple $Y_{i,t}$, one has a vector of P potentially explanatory factors, called features in the following. Let us relabel all the couples (i,t) by the index $n \in \{1, \dots, N\}$. The classification problem consists in explaining Y_n by a vector X_n with P components, or equivalently, to explain the vector $Y \in \{0, 1\}^N$ from the lines of matrix $X \in \mathbb{R}^{N \times P}$. Y is called the target and X the feature matrix. The problem is then to train a machine learning method to learn the mapping between the lines of X and the components of vector Y . Once the training is complete, such a model takes as input a vector of features and outputs the probability that these features correspond to one class (in a two-class problem).

The state of the art is Gradient Boosting models (Friedman, 2001), as shown for instance in Shwartz-Ziv and Armon (2021). The spirit of gradient boosting consists in using a sequence of weak learners (wrong models) that iteratively correct the mistakes of the previous ones, which eventually yields a strong learner (good model). We use here decision trees as weak learners. Different implementations of the Gradient Boosted Decision Trees method exist, e.g. XGBoost (Chen and Guestrin, 2016), LightGBM (Ke et al., 2017), CatBoost (Prokhorenkova et al., 2017). We use here LightGBM. The advantage of such methods with respect to logistic regression is that they are able to learn more generic, non-linear, functional forms and have accordingly superior performance.

The models are trained to minimize the cross-entropy (cost function), also known as LogLoss, defined as

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N y_i \log(p_i) + (1 - y_i) \log(1 - p_i),$$

where p_i is the model probability that sample i belongs in category 1 and $y_i \in \{0, 1\}$ is the true class (which selects the suitable term of the sum for each i). This type of loss implicitly assumes that both true classes appear with roughly the same frequency in the training set, which is the case with 51.7% of samples belonging to class 1 and 48.3% to class 0.

4.2 Training features

The Refinitiv ESG dataset contains several levels of granularity. We choose to train our models with the 10 pillar scores described in Appendix A (Resource Use, Emissions, Innovation, Workforce, Human Rights, Community, Product Responsibility, Management, Shareholders, CSR Strategy) and the aggregated Controversy score. This level of granularity is a good compromise.

We add five non-ESG features (market capitalization, country of incorporation and TRBC Sector at levels 1, 2 and 3). These features provide the benchmark features needed to settle the question of the additional information provided by ESG features .

4.3 Target computation

We compute the coefficients of the regression defined in Eq. 2 with monthly factors available online at Fama and French (2021) and monthly price returns over periods of 5 civil years. For instance, the regression coefficients used to compute the 2017 target possibly explained by 2017 ESG features are computed with historical data ranging from 2013 to 2017. We then compute targets over the year corresponding to the year of the ESG features: as we are in an explanatory setting, we want to explain the return of a company for a specific year by the ESG profile of this company during the same year.

4.4 Cross-validation and hyperparameter tuning in an increasingly good data universe

The usual strategy of a single data split into causal consecutive train, validation and test data sets may not be fully appropriate for the currently available ESG features. This is because the amount of data grows from a very low baseline both quantity–

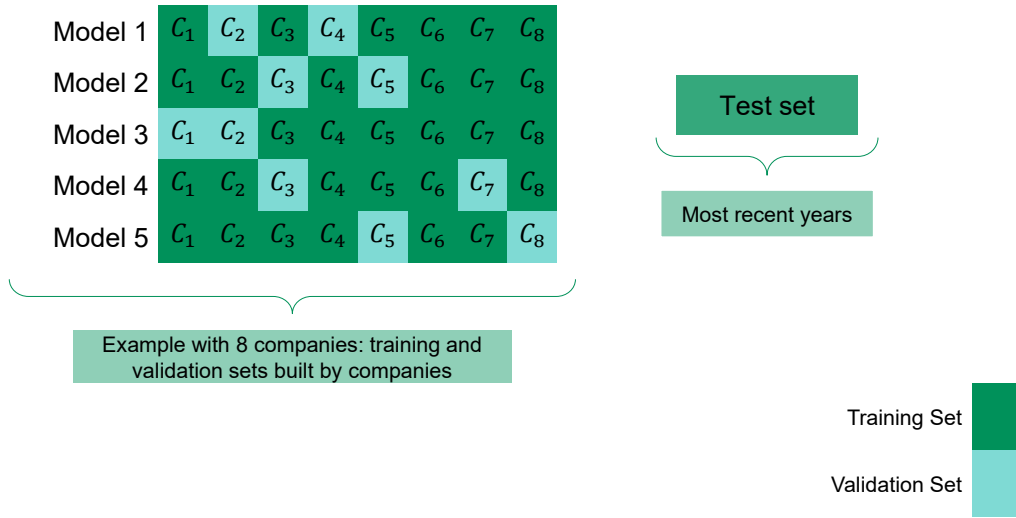


Figure 5: Company-wise cross-validation: the validation sets consists of randomly selected companies, which allows training to account for most of the most recent data.

and quality-wise, that was not exploitable, to an amount that more likely is. Thus, not only the data is non-stationary, but its reliability and quality keeps increasing. As a consequence, the cross-validation time splitting schemes known to work well in the context of non-stationary time series (Bergmeir and Benítez, 2012) may be improved on.

For this reason, we experiment with K -fold company-wise cross-validation where 75% of companies are randomly assigned to the train and the remaining 25% ones to the validation set (see Fig. 5). In other words, there are K different (train-validation) sets. For each of the K train sets, we train 180 models, varying 12 hyperparameters of the LightGBM (maximum tree depth, learning rate, etc.) and pool the five best ones according to model performance in the respective validation sets. In this way, models are trained with most of the most recent (hence, more relevant) data, while validating the model also with the most recent and best data. If the dependencies completely change every year, this validation scheme is bound to fail. As we shall see, this is not the case. We take $K = 5$.

In addition, we use expanding (train+validation)-test windows, using the last year as the test window, which allows us to perform a time-wise analysis of the performance of the models. Because data is insufficient before 2015, we have five different periods: the first test year is 2016 and the last one is 2020. We thus train and validate $K \times 5 = 25$ models.

For each testing period, we will compare the performance of the company-wise 5-fold random splits with that of the standard temporal split (75% train / 25% validation).


5 Results

We investigate here the results of the standard temporal split and the 5-fold company-wise split for a target computed using the CAPM model, as described in 4.1. Models trained using the Fama-French 3-factor model lead to less clear-cut performance; their results are relegated to Appendix B.


We first assess the quality of the models according to the cross-entropy loss, using their direct probability outputs. We also assess the end result, i.e., the predicted class. As it is usual, we map the output, a probability p_i , to classes 0 and 1 with respect to a 0.5 threshold. This allows us to compute the balanced accuracy, defined as the average of the sensitivity and the specificity. Sensitivity equals the ratio of true positives to the number of positive samples. Specificity is the ratio of the true negatives to the number of negative samples. An advantage of the balanced accuracy over classical accuracy is that balanced accuracy accounts for class unbalance in the test set. By definition, it assigns a score of 0.5 if the model did not learn anything significant.

We check that the performance of the models in the test sets bear some relationships with their performance in the validation sets. More precisely, for each (train+validation)-test period, we investigate the dependence between the cross-entropy losses in the validation and test sets for the best models trained during the hyperparameters random search, which makes it possible to characterize the training quality year by year. A significantly positive relationship shows that these models did learn persistent relationships, i.e., something useful. Mathematically, we assess the relationship $\mathcal{L}_m^{\text{test}}$ versus $\mathcal{L}_m^{\text{validation}}$ for each model m ranking for the 100 models with the best validation cross-entropy losses, for each of the five sets of (train+validation)-test sets. Figures 6a to 6e display these relationships for the company-wise cross-validation scheme and adds a linear fit. Figures of the same type for the standard time splitting scheme can be found in Appendix C. Generally, both test and validation cross-entropy losses are positively correlated, except for 2016. We believe that this comes from the fact that ESG data were of insufficient quality before that date. Year 2020 is also special: in addition to the coronavirus crisis, the data for 2020 was obtained at the beginning of 2021 when not all companies had ESG ratings, leading to a smaller dataset and a (mostly likely) biased test set.


We compute the Pearson correlation, the R^2 of the linear fit, Kendall tau and its p-value for the standard temporal split and the 5-fold company-wise split, which




(a) 2016




(b) 2017



(c) 2018



(d) 2019



(e) 2020

Figure 6: Company-wise cross-validation: test set cross-entropy versus validation cross-entropy of the 100 best models of the random hyperparameters search.

Company-wise 5-fold cross-validation				
Year	Pearson correlation	R^2	Kendall tau	p-value of Kendall tau
2016	-0.54	0.29	-0.36	$8.0e^{-8}$
2017	0.14	0.021	0.12	$6.7e^{-2}$
2018	0.47	0.22	0.30	$1.1e^{-5}$
2019	0.73	0.54	0.58	$1.5e^{-17}$
2020	0.27	0.071	0.19	$5.4e^{-3}$
Standard temporal split				
Year	Pearson correlation	R^2	Kendall tau	p-value of Kendall tau
2016	-0.43	0.18	-0.29	$1.6e^{-5}$
2017	0.46	0.21	0.33	$9.2e^{-7}$
2018	0.46	0.21	0.34	$7.7e^{-7}$
2019	0.47	0.22	0.33	$1.3e^{-6}$
2020	0.47	0.22	0.39	$7.6e^{-9}$

Table 1: Dependence measures between the cross-entropies (prediction error) in the validation and test sets, for the 100 best models of the random hyperparameters search.

are reported in Tab. 1. This latter allows us to compare the respective advantages and disadvantages of each validation strategy. All the dependence measures increase much from 2017 to 2019 for company-wise splits. The case of temporal split shows the limitations of this approach: the performance measures are roughly constant, which is consistent with the fact that adding one year of data to the train+validation dataset does not lead to much change. Display of the relationship $\mathcal{L}_m^{\text{test}}$ versus $\mathcal{L}_m^{\text{validation}}$ for the standard temporal model can be found in Appendix C, in Fig. 16.

Our second and most important aim is to establish that ESG data contains additional useful and exploitable information. To this end, for each train period defined above, we train a model with both ESG and benchmark features and another one with the benchmark features only. We assess both the absolute performance measures of the models and the amount of additional information brought by ESG features by computing the difference of performance measures in the test sets.

The company-wise splits make it easy to compute error bars on various measures: instead of training $K = 5$ models, we train 100 of them and then compute the median performance on 100 random subsets of size $K = 5$ among these 100 models. Table 2 provides results on absolute performance of the models for each test period, for both the company-wise and the standard temporal splits. Both splitting methods have a

Company-wise 5-fold cross-validation				
	Only Benchmark features		Benchmark and ESG features	
Year	Balanced Accuracy	Cross-entropy loss	Balanced Accuracy	Cross-entropy loss
2016	52.6	70.6	51.2	72.8
2017	57.4	69.2	56.9	69.6
2018	57.5	68.1	57.9	68.2
2019	65.6	63.1	67.9	62.7
2020	59.6	69.3	61.9	67.4
Standard temporal split				
	Only Benchmark features		Benchmark and ESG features	
Year	Balanced Accuracy	Cross-entropy loss	Balanced Accuracy	Cross-entropy loss
2016	53.2	68.8	51.8	70.3
2017	56.1	68.2	57.7	68.0
2018	56.2	67.5	58.1	67.4
2019	64.3	64.5	66.4	63.8
2020	58.5	70.5	61.0	69.6

Table 2: Performance measures in percent on the test set for both types of validation splits. The numbers for the company-wise splits are the median values of the performance of 100 random samplings of 5 models among 100 random company-wise validation splits.

clearly decreasing cross-entropy (a proxy for prediction error) as a function of time, except for 2020, which shows once again the special nature of this year in our dataset. This shows that the relevance of ESG features in price return formation increases as a function of time. Balanced accuracy displays a similar improvement before 2020. However, this time, company-wise splits yields increasingly better than temporal splits, which we believe is an encouraging sign of its ability to better leverage the latest and best data.

Figure 7 displays the time evolution of the cross-entropy and the balanced accuracy in the test sets. The boxplots are computed for the company-wise splits from the 100 associated predictions; the orange lines are the median of these performance measures, the rectangle delimits the first and third quartiles and extreme limits are situated before the first quartile minus 1.5 time the interquartile range and after the third quartile plus 1.5 time the interquartile range. Any point outside of this range is considered an outlier.

Company-wise 5-fold cross-validation outperforms the standard time splitting scheme, which supports our claim that the not fully mature nature of ESG data can be partly alleviated by a suitable validation scheme.

Figure 8 shows the difference of performance between the models trained on ESG and benchmark features and the models trained only on benchmark features, for the company-wise 5-fold cross-validation. ESG features contain more relevant

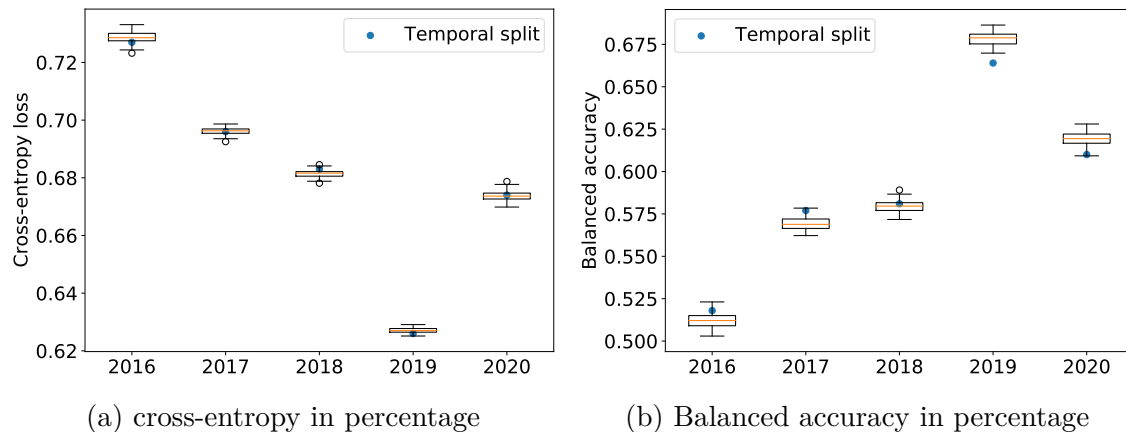


Figure 7: Performance measures on the test sets of the two train and validation schemes. The boxplots show the performance of 100 random samplings of 5 models among 100 random company-wise validation splits.

information as time goes on. Two explanations spring to mind: long positions are more and more driven by ESG-conscious investors, or the quality of data increases as a function of time, which makes the relevance of ESG scores more apparent.

6 Interpretability

We provide now a breakdown of the impact of the different ESG features on the predicted probability of having positive idiosyncratic returns in the CAPM model. Because of the superior performance of the company-wise K -fold cross-validation, we use this method in the following.

6.1 Shapley values

Shapley values, first introduced in the context of game theory (Shapley, 1953), provide a way to characterize how each feature contributes to the formation of the final predictions. Shapley values and their uses in the context of machine learning are well described in Molnar (2020).

The Shapley value of a feature can be obtained by averaging the difference of prediction between all the combinations of features containing and not containing the said feature. For each sample in our dataset, each feature possesses its own Shapley value representing the contribution of this feature to the prediction for this particular sample. Shapley values have very interesting properties, one of them being the efficiency property. If we note $\phi_{j,i}$ the Shapley value of feature j for a sample x_i and $\hat{f}(x_i)$ the prediction for the sample x_i , Shapley values must add up




Figure 8: Performance measures in comparison to benchmark, for the company-wise 5-fold cross-validation.

to the difference between the prediction for the sample x_i and the average of all predictions $E_X(\hat{f}(X))$ and then follow the following formula:

$$\sum_{j=1}^p \phi_j = \hat{f}(x) - E_X(\hat{f}(X)) \quad (3)$$

The dummy property also states that the Shapley value of a feature which does not change the prediction, whatever combinations of features it is added to, should have a Shapley value of 0.

Shapley values computation is quite time and memory intensive. Lundberg and Lee (2017) and later Lundberg et al. (2018) proposed a fast implementation of an algorithm called TreeSHAP, which allows to approximate Shapley values for trees models like the LightGBM, which we use in the following and refer to as SHAP values.

Let us just note that, as we are using a LightGBM model in classification, the prediction is not directly the probability of belonging to the class 1, but rather the logit associated with this probability. Probability is an increasing function of the logit and thus SHAP values obtained for the logit can easily be transformed for the probability. Indeed, for a sample x_i , the predicted probability of belonging to class 1 p_i is linked to the logit $logit_i$ according to :

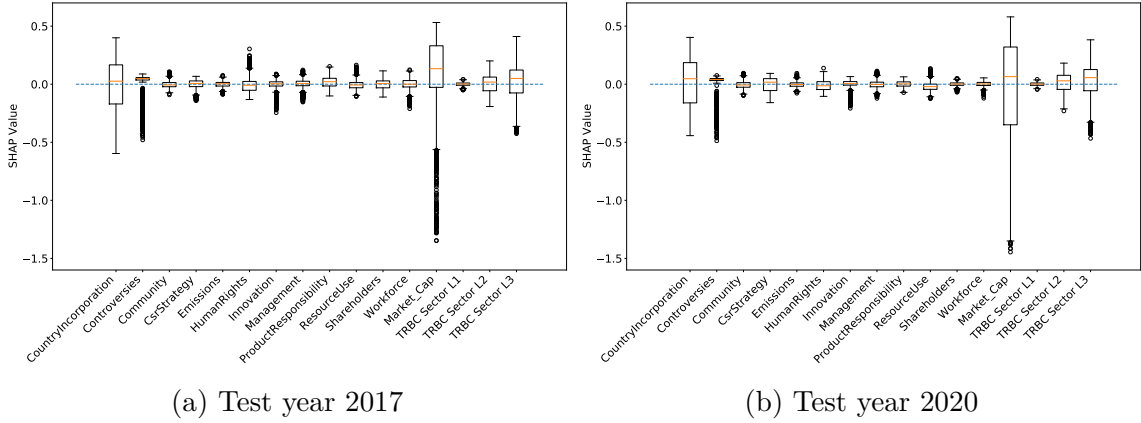


Figure 9: SHAP values distribution.

$$p_i = \frac{1}{1 + e^{-\text{logit}_i}} \quad (4)$$

6.1.1 Evolution of ESG features contribution from 2017 to 2020

In Fig. 9, we plot the distribution of SHAP values for each feature and for all test samples for models trained from 2002 to 2016 (9a) and trained from 2002 to 2019 (9b). The first teaching of this plot is that the contribution of ESG features to the predicted probability of having a positive return has not dramatically increased with the additional, more recent and more complete data. Benchmark features are the ones which have the biggest impact on the prediction. However, we observe an important number of outliers for some SHAP values associated with some features, demonstrating that these ESG features have more impact on the prediction for these particular samples. It would be interesting to study these outliers to understand more why ESG features are more important in explaining price returns for some samples than others.

For instance, we observe in Fig. 10 the score distributions for the outliers of the Controversy SHAP values. All of these scores are below 0.9, suggesting that the Controversy score is more informative when a company has indeed suffered controversies during the year and was then not able to reach a score of 1. Observing outliers of SHAP values and their associated scores, we can make the hypothesis that ESG features are important and have a strong impact in the explanations of past returns if their score is extreme. This would mean that ESG information would lie in extreme scores, more standard scores bringing much less information. Checking this hypothesis is beyond the scope of this work and is left for future investigations.




Figure 10: Distribution of data for lowest outliers of SHAP values of 2020, for Controversy score.

6.2 Partial dependence plots: marginal effect of ESG features

6.2.1 Definition

A partial dependence plot shows the marginal effect of features on the prediction made by the model. It is a way of understanding the links the model made from features to the target, and that it had understood from the data. It also shows if this relation is linear or not, monotonic or not, . . . Partial dependence plots were first introduced in Friedman (2001) and are also well described in Molnar (2020). Briefly, a partial dependence plot for a feature of interest is obtained by marginalizing the predicted output over the values of all other input features. This marginalization is done by calculating averages in the training data, using a Monte-Carlo method, with a fixed value for the features of interest.

An important limitation of a partial dependence plot is that their methodology of construction assumes independence between the features, which does not seem to be the case for ESG features. This limitation is neglected here. All partial dependence plots are made with the most recent model, trained with data from 2002 to 2019, on a subsample of recent ESG data.

6.2.2 Marginal effect of the ESG features

Using partial dependence plots, we first compute the marginal effect of each ESG feature on the probability of having a positive return during the year of publication of the ESG features (Fig. 11). Figure 12 reports the sector by sector probability of having a positive predicted return.

Figure 11 shows that ESG features are mostly not related in a monotonic way with the probability of having a positive return. A clear exception would be the

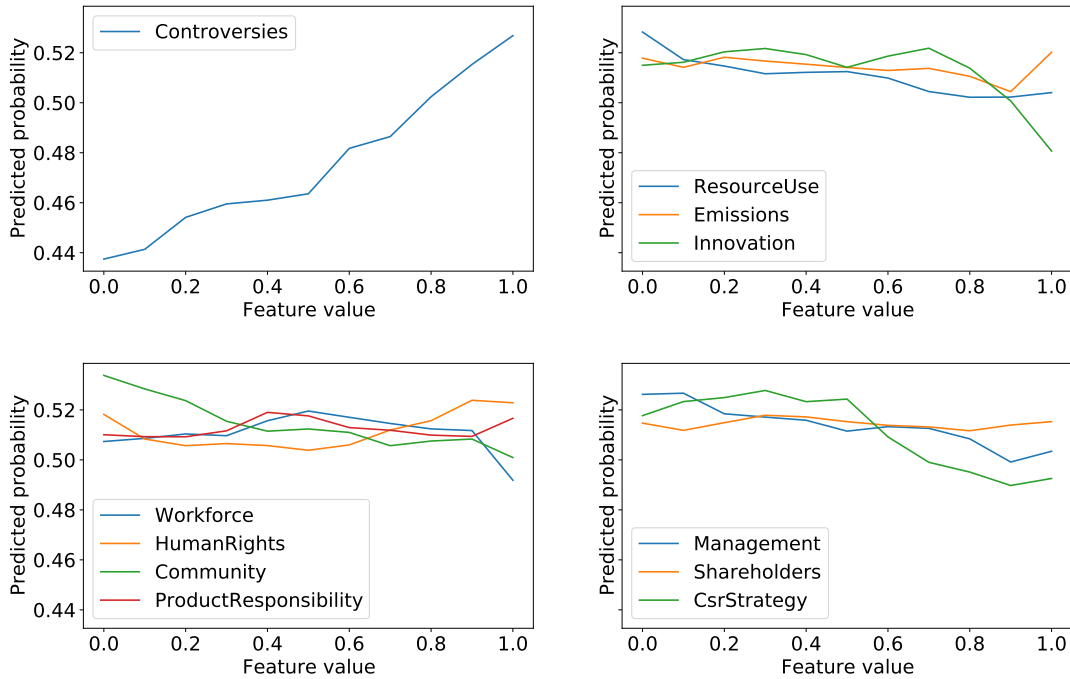


Figure 11: Marginal effect of each ESG feature on the predicted probability of having a positive return.

Controversy score, on the top left, which shows a strong monotonic relation which strongly implies that being subject to controversies during a year leads to a lower probability of having a positive return. For the 10 pillar scores, one sees a much weaker dependence. For example, the probability of positive price return increases by around 1% when the Product Responsibility and Shareholders scores increase from 0 to 1. Still, a trend is present for most of these ESG features: partial dependence plots for features such as Resource Use, Innovation, Community or Management seems to be decreasing, suggesting that obtaining better ESG scores and practices comes at the price of a slightly degraded financial performance.

6.2.3 Marginal effect of the ESG features sector by sector: materiality matrices

Adding the section dimension to partial dependence plots yields so-called materiality matrices. In our setting, it is a table whose rows represent ESG features and whose columns are economic sectors. A cell of this matrix shows, in percentage, by how much the probability of having a positive return is increased by going from a low score (between 0 and 0.2) to a high one (by 0.8 to 1). This quantity is easily obtained using partial dependence plots: for a specific selected economic sector, we can plot




Figure 12: Marginal effect of the sector (TRBC Sector L1) feature on the predicted probability of having a positive return.

the evolution of the predicted probability against the feature value. Making the strong hypothesis of a monotonic and close to linear relationship, we can compute the value in the cell as the slope of the trend line of the precedent plot.

The obtained materiality matrix is presented in Fig. 13. All the TRBC Level 1 sectors are included. Results for Academic & Educational Services should be handled with care as they are not based on as many samples as the ones for other sectors, as shown in Fig. 4. Some ESG scores have some strong impact on the probability of having positive returns. The Controversy score especially has a similar impact for all sectors: not suffering controversies during the year increases the probability of having a positive return. On the contrary, the CSR Strategy row shows that working towards the integration of social and environmental dimensions into the day-to-day decision-making processes, in addition to economic and financial ones, leads to a loss of financial performance. It is also the case for Resource Use, Environmental Innovation, Community, and Management scores, each with a different magnitude.

Furthermore, we bucket the companies which serve to build this materiality matrix by market capitalization. We choose three buckets, with small market capitalization being below 2 billion euros, mid ones between 2 and 10 billion euros and large ones above 10 billion euros, which correspond to the ones Refinitiv use when calculating the Controversy score. The three obtained materiality matrices are presented in Fig. 14. The marginal effect of the Controversy score remains the same, even if it is slightly smaller for the small caps. However, companies with a

large market capitalization benefit from a better impact of ESG: for some features, working toward better ESG scores can preserve or even boost financial performance, whereas it would be the opposite for small caps. For instance, large caps companies have an average materiality of 0.8 for the Resource Use score and 1.5 for the Emissions scores, whereas small caps ones have respectively average scores of -4.6 and -1.1, denoting a clear difference.

To obtain a statistically meaningful interpretation of these results, we need to account for the fact that to each cell corresponds coefficients of a linear fit with associated p-values, i.e., one makes one null hypothesis per cell. We thus need to use multiple hypothesis correction to check globally which cells are statistically significant results. Here, we choose to control the False Discovery Rate with the Benjamini-Hochberg procedure (Benjamini and Hochberg, 1995). We set the FDR to 5%, which means that there are only about 3 false discoveries in each of the reported tables.

7 Conclusion

While ESG data are not yet fully mature and lack long enough quality records to be amenable to easy conclusions, powerful machine learning and validation techniques make it already possible to show that they do influence yearly price returns, and increasingly so. We stress that our results are specific to the Refinitiv ESG dataset, for the European market. Breaking down their influence sector by sector, subscore-wise and according to market capitalization clearly demonstrates that an average approach will fail to be informative. We found in particular that the relationship between controversies and price return is the most robust one. The average influence of all the other ESG scores depends much on the capitalization of a company: strikingly, most of the statistically significantly influential ESG scores weigh negatively on price returns of small or mid-size companies. Large-capitalization companies on the other hand have significantly advantageous ESG score types.

While this work proposes a methodology suitable to an explanatory setting, our methodology can be used to explore the predictive power of ESG data provided that data revision is under control.

While this work focuses on explaining the idiosyncratic part of price returns derived from the CAPM model, those derived from the Fama-French 3-factor model lead to results that are less clear cut for the time being but less and less so as time goes on: the correlations between the errors in the validation and test sets increased in 2018 and 2019 and thus ESG data seems to become more informative. Future work will focus on the study of the full 2020 and 2021 years to check if this is confirmed.

Future work will also include studying outliers of the SHAP values distribution






Figure 13: Materiality matrix: marginal effects of the combination ESG feature/Sector feature on the predicted probability of having a positive return. Blank cells are those which were not found statistically significant by the Benjamini-Hochberg procedure.



(a) Small market capitalization (<2B) (b) Mid market capitalization (>2B, <10B)



(c) Large market capitalization (>10B)

Figure 14: Materiality matrices: marginal effects of the combination ESG feature/Sector feature on the predicted probability of having a positive return, bucketed by market capitalization. Blank cells are those which were not found statistically significant by the BenjaminiHochberg procedure.

and verifying the hypothesis that extreme scores in the ESG field are more informative. In addition, the link between ESG and equity returns is complete only if the systematic and idiosyncratic aspects of risks and returns are studied together (Giese and Lee, 2019): indeed, it may be that having better ESG scores not only decreases price returns but also risk.

References

- Fabio Alessandrini and Eric Jondeau. ESG investing: From sin stocks to smart beta. *The Journal of Portfolio Management*, 46(3):75–94, 2020.
- Mark Anson, Deborah Spalding, Kristofer Kwait, and John Delano. The sustainability conundrum. *The Journal of Portfolio Management*, 46(4):124–138, 2020.
- Steven Bacon and Arnfried Ossen. Smart ESG integration: Factoring in sustainability. *RobecoSam AG*, 2015.
- Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)*, 57(1):289–300, 1995.
- Christoph Bergmeir and José M Benítez. On the use of cross-validation for time series predictor evaluation. *Information Sciences*, 191:192–213, 2012.
- André Breedt, Stefano Ciliberti, Stanislao Gualdi, and Philip Seager. Is ESG an equity factor or just an investment guide? *The Journal of Investing*, 28(2):32–42, 2019.
- Michael Cappucci. The ESG integration paradox. *Journal of Applied Corporate Finance*, 30(2):22–28, 2018.
- Mike Chen and George Mussalli. An integrated approach to quantitative ESG investing. *The Journal of Portfolio Management*, 46(3):65–74, 2020.
- Tianqi Chen and Carlos Guestrin. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794, 2016.
- Bradford Cornell and Aswath Damodaran. Valuing ESG: Doing good or sounding good? *NYU Stern School of Business*, 2020.
- Carmine De Franco, Christophe Geissler, Vincent Margot, and Bruno Monnier. ESG investments: Filtering versus machine learning approaches. *arXiv preprint arXiv:2002.07477*, 2020.
- Eugene F Fama and Kenneth R French. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33:3–56, 1993.
- Eugene F Fama and Kenneth R French. Fama and French portfolios and factors data, 2021. URL https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

- Gunnar Friede, Timo Busch, and Alexander Bassen. ESG and financial performance: aggregated evidence from more than 2000 empirical studies. *Journal of Sustainable Finance & Investment*, 5(4):210–233, 2015.
- Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232, 2001.
- Guido Giese and Linda-Eling Lee. Weighing the evidence: ESG and equity returns. *MSCI Research Insight*, 2019.
- Douglas M Grim and Daniel B Berkowitz. ESG, SRI, and impact investing: A primer for decision-making. *The Journal of Impact and ESG Investing*, 1(1):47–65, 2020.
- Tian Guo, Nicolas Jamet, Valentin Betrix, Louis-Alexandre Piquet, and Emmanuel Hauptmann. ESG2risk: A deep learning framework from ESG news to stock volatility prediction. *arXiv preprint arXiv:2005.02527*, 2020.
- Adolfo Hilario-Caballero, Ana Garcia-Bernabeu, Jose V Salcedo, and Marisa Vercher. Tri-criterion model for constructing low-carbon mutual fund portfolios: a preference-based multi-objective genetic algorithm approach. *International Journal of Environmental Research and Public Health*, 17(17):6324, 2020.
- Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. LightGBM: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30:3146–3154, 2017.
- Marlene H Konqui, François Millet, and Serge Darolles. Why using ESG helps you build better porfolios, 2019.
- Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Proceedings of the 31st international conference on neural information processing systems*, pages 4768–4777, 2017.
- Scott M Lundberg, Gabriel G Erion, and Su-In Lee. Consistent individualized feature attribution for tree ensembles. *arXiv preprint arXiv:1802.03888*, 2018.
- Christoph Molnar. *Interpretable machine learning*. Lulu. com, 2020.
- Jan-Carl Plagge and Douglas M Grim. Have investors paid a performance price? examining the behavior of ESG equity funds. *The Journal of Portfolio Management*, 46(3):123–140, 2020.
- Liudmila Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna V Dorogush, and Andrey Gulin. CatBoost: unbiased boosting with categorical features. *arXiv preprint arXiv:1706.09516*, 2017.

- Refinitiv. Environmental, social and governance (ESG) scores, 2021.
- Anthony A Renshaw. ESGs evolving performance: First, do no harm. Frankfurt: Axioma, 2018.
- Thierry Roncalli and Vincent Mortier. The alpha and beta of ESG investing, 2019.
- David Schofield, Adam Craig, and Richard Yasenchak. What to look for on the road to ESG. *Intech - Janus Henderson*, 2019.
- Llyod S Shapley. A value for n-person games. *Contributions to the Theory of Games*, 2(28):307–317, 1953.
- Ravid Shwartz-Ziv and Amitai Armon. Tabular data: Deep learning is not all you need. *arXiv preprint arXiv:2106.03253*, 2021.
- Robert Tibshirani. Regression shrinkage and selection via the LASSO. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- Blaine Townsend. From SRI to ESG: The origins of socially responsible and sustainable investing. *The Journal of Impact and ESG Investing*, 1(1):10–25, 2020.

A 10 Pillar Scores

A.1 Environmental scores

- Resource Use: Reduce the use of natural resources and find more eco-efficient solutions by improving supply chain management.
- Emissions: Commitment and effectiveness towards reducing environmental emissions in the production and operational processes.
- Innovation: Reduce the environmental costs for customers, and thereby creating new market opportunities through new environmental technologies and processes or eco-designed products.

A.2 Social scores

- Workforce: Job satisfaction, healthy and safe workplace, maintaining diversity and equal opportunities, development opportunities for workforce.
- Human Rights: Respecting the fundamental human rights conventions.
- Community: Commitment towards being a good citizen, protecting public health and respecting business ethics.
- Product Responsibility: Producing quality goods and services integrating the customer's health and safety, integrity and data privacy.

A.3 Governance scores

- Management: Commitment and effectiveness towards following best practice corporate governance principles. Composition, remuneration, transparency of the board.
- Shareholders: Equal treatment of shareholders, use of anti-takeover devices.
- CSR Strategy: Integration of social and environmental dimensions into the day-to-day decision-making processes, in addition to economic and financial ones.

Company-wise 5-fold cross-validation				
Year	Pearson correlation	R^2	Kendall tau	p-value of Kendall tau
2016	-0.23	0.052	-0.14	4.4^{-2}
2017	-0.054	0.0030	0.010	8.8^{-1}
2018	0.29	0.085	0.19	4.2^{-3}
2019	0.67	0.44	0.49	7.1^{-13}
2020	0.053	0.0028	0.017	8.0^{-1}


Table 3: Dependence measures between the cross-entropy losses in the validation and test sets, for the 100 best models of the random hyperparameters search, for a target computed using the Fama-French 3-factor model.

Company-wise 5-fold cross-validation				
Year	Only Benchmark features		Benchmark and ESG features	
	Balanced Accuracy	Cross-entropy loss	Balanced Accuracy	Cross-entropy loss
2016	57.9	65.8	56.0	66.7
2017	55.0	70.6	55.2	71.6
2018	56.0	70.4	56.0	71.1
2019	62.4	64.6	64.7	64.1
2020	56.1	72.2	55.3	71.3


Table 4: Performance measures in percent on the test set, for a target computed using the Fama-French 3-factor model.

B Results with the target derived from the Fama-French 3-factor model


The following results were obtained with a target derived from the Fama-French 3-factor model as exposed in 4.1. This target was not selected as results were not as good as those obtained with the target derived from the CAPM model. How to interpret results with this target, especially in terms of materiality matrices, was also less clear. For the interested reader, we present our results, using a 5-fold company-wise splitting strategy, in Tab. 3 and 4. Display of the relationship between $\mathcal{L}_m^{\text{test}}$ versus $\mathcal{L}_m^{\text{validation}}$ for each model m ranking in the top 100 validation cross-entropy losses are shown in Fig. 15.




(a) Slope: -0.42 - R^2 : 0.05 -
Test year 2016




(b) Slope: -0.12 - R^2 : 0.00 -
Test year 2017



(c) Slope: 1.00 - R^2 : 0.08 -
Test year 2018




(d) Slope: 1.93 - R^2 : 0.44 -
Test year 2019




(e) Slope: 0.14 - R^2 : 0.00 -
Test year 2020


Figure 15: Company-wise cross-validation: test set cross-entropy versus validation cross-entropy of the 100 best models of the random hyperparameters search, for a target computed using the Fama-French 3-factor model.




(a) 2016




(b) 2017



(c) 2018



(d) 2019



(e) 2020

Figure 16: Standard temporal split: test set cross-entropy versus validation cross-entropy of the 100 best models of the random hyperparameters search.

C Relationship between validation and test cross-entropy losses for the temporal train/validation scheme

We assess the relationship $\mathcal{L}_m^{\text{test}}$ versus $\mathcal{L}_m^{\text{validation}}$ for each model m ranking for the 100 models with the best validation cross-entropy losses, for each of the five sets of (train+validation)-test sets. Figures 16a to 16e display these relationships for the standard time splitting scheme.