

ESG Investments: Filtering versus Machine Learning Approaches

Vincent Margot¹, Christophe Geissler¹, Carmine de Franco², Bruno Monnier²

¹ Advestis, France

² Ossiam, France

Correspondence: Vincent Margot, Advestis, France.

Received: December 7, 2020

Accepted: January 22, 2021

Available online: February 8, 2021

doi:10.11114/aef.v8i2.5097

URL: <https://doi.org/10.11114/aef.v8i2.5097>

Abstract

We designed a machine learning algorithm that identifies patterns between ESG profiles and financial performances for companies in a large investment universe. The algorithm consists of regularly updated sets of rules that map regions into the high-dimensional space of ESG features to excess return predictions. The final aggregated predictions are transformed into scores which allow us to design simple strategies that screen the investment universe for stocks with positive scores. By linking the ESG features with financial performances in a non-linear way, our strategy based upon our machine learning algorithm turns out to be an efficient stock picking tool, which outperforms classic strategies that screen stocks according to their ESG ratings, as the popular best-in-class approach. Our paper brings new ideas in the growing field of financial literature that investigates the links between ESG behavior and the economy. We show indeed that there is clearly some form of alpha in the ESG profile of a company, but that this alpha can be accessed only with powerful, non-linear techniques such as machine learning.

Keywords: best-in-class approach, ESG, machine learning, portfolio construction, sustainable investments

1. Introduction

The relationship between corporate social (CSP) and financial performances (CFP) is fairly an old theme in the economic research. In its earlier stages, it has met quite deep skepticism and critics: Nobel prize-winning economist Milton Friedman wrote in the New York Times Magazine, back in the 1970', that "... *there is one and only one social responsibility of business - to use its resources and engage in activities designed to increase its profits so long as it stays within the rules of the game, which is to say, engages in open and free competition without deception or fraud...*" (Friedman 1970).

But the number of studies that highlights positive, or at least non-negative, relationship between social and financial performances has grown significantly since then, probably beginning with the initial work by Bragdon and Marlin (1972) on the link between environmental virtue and financial performance. Fifty years later the context has completely changed. The number of proponents of social and, more broadly, ESG integration in both corporate management and investors' choices has grown exponentially. And so has the number of financial products, funds and ETFs, that offer ESG versions of a large panel of investment strategies (mainly on equity and bonds).

The current approach now seems completely opposed to Friedman's one, and the most recent empirical literature highlights the link between ESG performance and alpha (Chong & Phillips, 2016; Giese, Ossen, & Bacon, 2016; Zoltan, Kassam, & Lee, 2016). Nonetheless, the question regarding the relationship between CSP and CFP remains unanswered. Reviews of published paper (meta-analysis) highlight that the majority of empirical studies published on this theme reports non-negative or small positive relationship between CSP and CFP (see for example Orlitzky, Schmidt, & Rynes, 2003; Allouche & Laroche, 2005; Wu, 2006; Van Beurden & Gössling, 2008; Margolis, Elfenbein, & Walsh, 2009; Friede, Bush, & Bassen, 2015). Other researchers take a more optimistic view and report significant relationship between CSP and CFP (Peiris & Evans, 2010; Filbeck, Holzhauser, & Zhao, 2014; Indrani & Clayman, 2015) or at least that CSP is not detrimental to CFP as long as one manages to build the portfolio with care, even if there is no clear value added in ESG integration (Kurtz & Di Bartolomeo, 2011).

Although we do not share the very optimistic, and mostly overstated, enthusiasm about the direct relationship between ESG and financial performance, we do believe that there is a strong relationship between ESG and sustainability of corporate business. Therefore, ESG has an impact on financial performances and risks, but this does not come linearly.

We welcome the efforts that investors are undertaking to include ESG criteria into their portfolio choices, and we clearly hope that this will trigger economic and cultural changes in corporate management. At the same time, we remain skeptic in front of the far too flaunted capability of basic ESG ratings to act as an alpha generator in a portfolio.

It remains true however that the very large set of ESG data, reports and analysis can contain useful information related to the strengths and weaknesses of corporations. Unfortunately, ESG ratings are, by construction, a composite measure that dramatically reduces this rich set of information.

Our contribution to the growing literature on this topic is to show that, empirically, there is no value added in portfolios based on simple ESG screenings. Although it usually comes with no harm to the performance, we do not find any alpha in such approaches. However, by recognizing the intrinsic value of the large panel of ESG indicators that are aggregated to form the ESG ratings, we show that it is possible to extract value from them, which, in turns, translates into real alpha. By exploring large data sets of specific ESG indicators, we are able to identify those who really have an impact on corporate financial performances. In a simplified example, we can agree on the fact that for a company in the utility sector, most likely, the environmental performance can be a discriminating criterion for financial performance; at the same time, governance can play an important role if we compare a utility company in Europe with one in an emerging country. Similarly, direct carbon emissions for banks are probably not as relevant as the exposures of these banks, through loans, to highly polluting companies would be. In short, aggregate measures as ESG ratings lose valuable information contained in the ESG indicators, which therefore lower their predictive power.

Searching for interesting patterns between specific ESG indicators and financial performance for a large set of companies remains out of reach for the standard tools available to econometricians. This search takes place in a very high-dimensional space and is not oriented by some a priori on these ESG features. To deal with this complexity, we developed a machine learning algorithm that allows us to identify features and patterns that are relevant to explain the link between CSP and CFP. The algorithm maps the regions in our high-dimensional space of ESG features that have been consistently associated with outperformance or underperformance. In the econometric parlance, we look at those regions for which the conditional expectation of each stock's forward return is statistically positive (or negative), given that its relevant ESG features fall in these regions. We say that these relevant ESG features "activate" the region. By observing the ESG features we then obtain a significant signal regarding the future financial performance of the stock.

This identification is done with a set of rules that take the form of If-Then statements. The If statement identifies the region in the ESG space: in other words, the values that some ESG features have to take in order to activate the rule. The Then statement produces a prediction of the excess return, over the benchmark, that we can expect from a stock whose ESG features fall in that region. The final prediction is the aggregation of the predictions made by these rules and is transformed into a score. We therefore focus on the sign of the prediction of excess return rather than on its value. This usually makes the estimation more robust. The aggregation method mimics a panel of experts, each of which is expert on a very specific ESG feature (ex. Environment, Independence of the Board, ESG Reporting Verification, Employee Incidents, etc.) and makes a prediction given the ESG behavior of the company. When the aggregated prediction is close to zero, i.e. the panel of experts is split between optimistic and pessimist forecasters, the final prediction is set at 0. The algorithm is regularly trained over time so that it can react and readjust to the new observed data.

The algorithm is used to design a very simple strategy that screens the investment universe and selects all stocks with a positive score. The resulting portfolio is compared with a classic ESG best-in-class portfolio, which consists of all stocks in the investment universe whose ESG ratings are above a given threshold within their peer groups. Our empirical results show that the simple machine learning screened portfolio significantly outperforms the ESG best-in-class approach and the benchmark.

This is in line with the economic belief that ESG data is valuable information to assess financial performance, but also confirms that aggregated ESG ratings are not suited to distinguish between outperformers and underperformers over the long run. Even if perfect distinction is out of reach, our results clearly point out the fact that there is alpha in the granular ESG data, but the relation between ESG and financial performance is definitely not linear. Furthermore, the predictive power of the scores vanishes with time. We proved indeed that regularly training the algorithm over time, and producing up-to-date sets of rules, are key components of the superior performance of the machine learning when it comes to stocks screening.

2. Data

The analyses in this paper are carried out on portfolios based on the investment universe defined by the capitalization-weighted MSCI World Index USD, that consists of the largest capitalization listed in the US, Canada, Western Europe, Japan, Australia, New Zealand, Hong Kong and Singapore. Portfolios are calculated in USD and net dividends are reinvested in the portfolio itself. Stock prices and dividends are taken from Thomson Reuters/Datastream.

We reconstruct a proxy of the MSCI World Index by using end-of-month compositions as well as proxies for the US, Europe¹ and Asia Developed² benchmarks. We also consider sector portfolios derived from the MSCI World Index and the regional benchmarks by filtering on stocks that belong to the same sector: Consumer Staples (CS), Consumer Discretionary (CD), Energy (EN), Financials (FI), HealthCare (HC), Industrials (IN), Information Technology (IT), Materials (MA), Telecommunication Services (TL), Utilities (UT).

For each company in the investment universe, we collect ESG ratings from Sustainalytics³. An ESG rating is a comprehensive measure based on three pillars, Environment, Social and Governance, that assesses the strengths and weaknesses of a company along these three directions. The pillars are themselves based on a large set of specific indicators. For the purposes of this study, the composite ESG rating is the arithmetic average of the three ratings Environment (E), Social (S) and Governance (G), each of which is itself the combination of roughly 50 narrower indicators. Finally, for each company, we consider its relative Peer Group, which consists of all companies with a similar business, hence comparable from a sustainability point of view.

3. The Best-In-Class Approach

One of the most popular approaches to embed ESG criteria in the portfolio construction process is the so-called best-in-class approach. Given a threshold x , one excludes the stocks whose ESG ratings belong to the lowest x -quantile. The exclusion is usually carried across peer groups, i.e. groups of stocks with very similar characteristics. The reason behind this is twofold:

- Removing stocks with low ESG ratings within peer groups ensures that the final economic mesh of the filtered universe remains similar to the initial investment universe.
- ESG ratings have a structural, sector-driven bias which usually favor specific sectors (ex. IT or HealthCare sectors) while penalizing others (ex. Energy or Utilities). Given this bias, the filtering over peer groups makes comparisons of ESG ratings independent of the sectors.

For the purpose of this study, an ESG best-in-class portfolio derived from a capitalization weighted portfolio removes, within each peer group, the stocks whose ratings belong to the lowest x -quantile. The portfolio is finally scaled to sum up to one.

This approach, quite popular among investors, should not be thought of as a way to enhance performance. As Tables 1 - 4 in Annexes, ESG best-in-class filters applied to standard capitalization-weighted indexes do not bring outperformance.

Except for Europe and relatively low threshold levels, we find small but negative excess returns and negative information ratios for the ESG best-in-class portfolios over their benchmarks with almost unchanged risks. Although the approach does not create outperformance per se, it does not carry structural under-performance either. Optimistically, one could accept the fact that embedding ESG objectives in a portfolio does not significantly modify its risk/return profile.

Our findings are not in contradiction with the large literature that finds positive links between ESG and financial performance. But the consistency and durability over time of the ESG factor has been questioned since the very beginning. Aupperle, Carroll, & Hatfield (1985) find no significant relationship between social responsibility and corporate profitability, and similar results were obtained in Capelle-Blancard and Monjon (2012) and Humphrey and Tan (2014). Griffin and Mahon (1997) report that correlation between financial performance and social performance depends on the measure used to distinguish between high and low social performers.

Our results are more in line with Revelli and Viviani (2015) for which "... the consideration of corporate social responsibility in stock market portfolios is neither a weakness nor a strength compared with conventional investments...". It should be noted that many fund managers and institutional investors surveys report that ESG is mostly looked as a risk mitigation tool in the first place (Van Duuren, Plantinga, & Scholtens, 2016), and eventually as a performance driver at longer horizon. We share the optimistic view of Nobel prize-winning economist Robert Shiller for which both society and the financial community would find the use of socially responsible practices mutually beneficial (Shiller, 2013). At the same time, we also believe that short to mid term financial performance is at best lowly correlated to ESG ratings, at least for such broad investment universes as the MSCI World Index (which contains more than 1,600 companies). We can list several reasons for this:

¹ Stocks in the MSCI World Index domiciled in Austria, Belgium, Denmark, Finland, France, Germany, Ireland, Italy, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland and United Kingdom.

² Stocks in the MSCI World Index domiciled in Australia, Hong Kong, Japan, New Zealand and Singapore.

³ One of the largest provider of ESG ratings.

1. The investment universes are relatively large and the aggregated ESG ratings have too little a signal-to-noise ratio to allow for an efficient selection of outperforming stocks.
2. ESG ratings are global metrics that embrace environmental, social and governance criteria. As such, they may be too reductive and we may lose a significant amount of information from the single indicator to the aggregated scores.
3. Granularity is key: As an example, it is likely that companies in specific sectors (ex. Energy) react differently to changes in the environment score (E) compared to the social score (S).
4. In the search for a rational economic theory behind ESG, some argue that by divesting low ESG rated companies, investors raise their cost of capital and, in turns, the return these companies have to offer to attract new investors. As such, in the short run, they may show higher performances, but over time, the level of return they have to offer becomes unsustainable. Said otherwise, the action of divesting may take time to materialize in both investors' portfolios and low ESG rated companies (see for example Asness, 2017).
5. The period considered in this study spans from the earlier stages of the recovery in 2009 to March 2018. Therefore, we are considering key performance indicators over a period of strong equity market, characterized by high returns and historically lower levels of volatility. This market regime can potentially affect the overall strength of ESG filtered portfolios.

To illustrate item (3), we consider sector portfolios derived from the MSCI World Index and from the three regional benchmarks (US, Europe, Asia) and we apply both ESG and single pillar E, S and G, 30% best-in-class filtering. Tables 5 - 8 in Annexes collect the results. For the sake of simplicity, we only show annualized excess returns over the relative benchmark sector portfolios and information ratios.

Overall, it is not straightforward to detect clear patterns between excess returns and ESG metrics conditionally to the regional benchmarks. But we can definitely detect specific triplets sector/region/metric that produce significant positive excess returns. Clearly, integrating ESG criteria in the Utilities sector enhances in-sample performances. But the right metric to use clearly depends on the geography: in the World Developed (Table 5) the best excess return for the Utilities sector is achieved when one uses the Governance (G) score only at 0.82%; in the US (Table 6) it is better to look at the composite ESG ratings which achieves 0.63%. In Europe (Table 7) it is the Environment score (E) that obtains the best result with 3.25% while in the Asia (Table 8) it is, once again, the composite ESG rating that achieves the highest excess return at 3.07%.

More generally, there is no sector nor metric for which the excess return of the best-in-class filtered sector achieves positive excess return in all the regions. Similarly, there is no region nor sector for which all metrics produce positive excess returns. Finally, no sector achieves positive excess returns across all regions and metrics. In other words, finding performance drivers when integrating ESG criteria in a best-in-class fashion is out of reach.

From Tables 5 - 8, only 12 out of 40 sector/metric portfolios in the World Developed region turn out to have positive excess return, and half of them are obtained when one considers the Governance (G) score. In the US we find positive excess returns in 14 out of 40, with no clear indication on the best metric to use. We notice though that all the metrics seem to work in the Utilities sector.

In Europe, we count 25 out of 40 sector/metric pairs with positive excess returns. For 4 sectors (Consumer Discretionary, Materials, Telecommunication Services and Utilities) all metrics work accurately. In Asia, we have 16 out of 40 portfolios with positive excess returns with no clear patterns between sectors and metrics, except for the Energy sector for which all metrics produce positive excess returns, even if their magnitudes are relatively small.

In conclusion, our empirical findings confirm that simple ESG filtering does not bring extra performance. Overall, it rather behaves as a small drag. Given the short period we consider, and the market regime that equity markets have experienced since 2009, we share the view that ESG best-in-class integration is, most likely, neutral to financial performance. Nevertheless, our results highlight the fact that geographies and sectors do not react to ESG criteria in the same way. But finding interesting and statistically significant patterns between ratings, pillars, their underlying narrow indicators (features) and financial performances, for more than 150 indicators on more than 1,600 companies in the MSCI World Index, over roughly 10 year, is out of reach for both human and linear statistic tools.

Next section introduces other techniques that can overcome this complexity and exploit this huge set of data.

4. Machine Learning

In this section we introduce a deterministic, easily understandable machine-learning prediction algorithm, aimed at finding consistent and statistically significant patterns between ESG ratings and financial performances. The algorithm explores a high-dimensional dataset of ESG granular indicators for all the companies in our investment universe.

The goal of the algorithm, which falls in the category of supervised machine learning, is to predict the (conditional) excess return of each company over the benchmark, given the specific values taken by some of its ESG indicators (the features). Stated differently, the algorithm identifies regions in the high-dimensional space of ESG features that are statistically related to financial outperformance or underperformance. Features include raw and derived ESG indicators⁴, sector and country classifications, company's size and controversy indicators.

The regions are characterized by rules in the form *If-Then*, so that the algorithm finally consists of a set of such rules. The *If* statement is a list of conditions on the features $x_t \in \mathbb{X} = X_1 \times \dots \times X_d$, where X_i is the set of possible outcomes of the feature i and $d = 447$ is the total number of features⁵. Therefore, a rule defines a hyper-rectangle of \mathbb{X} . The *Then* statement is the prediction of the 3-month forward excess return conditionally to the *If* statement. Since the rules correspond to hyper-rectangles in the feature space, we finally obtain relatively simple and understandable regions. Furthermore, to avoid over-fitting, the algorithm only selects a finite number of such rules. At each time t , the predictions of each rule are aggregated into one prediction, \hat{y}_t , through convex combination.

The algorithm is calibrated (*trained*) on the training set and the rules are used out-of-sample. The learning process works at two independent levels:

- At the end of year $N + 1$ we train the algorithm on an expanded data-set of features and stock total returns which contains the data-set used at the end of year N augmented of all the new observed data (features and stock total returns) from the end of year N to the end of year $N + 1$. To initialize the algorithm, we train it over 3 year of data (from 2009 to 2012). By expanding the dataset, the algorithm is able to access new data and explore new patterns, so that it can strengthen or nuance some rules that were previously discovered.
- On a daily basis, the algorithm can update the weights used to aggregate each rule's prediction, by over-weighting rules with a good prediction rate and under-weighting the others. Therefore, following day predictions will benefit from the experience the algorithm is gaining on the rules and their predictive power. The weight of each rule can be viewed as a confidence index. Of course, this is possible because the algorithm is able to assess the goodness of its predictions by looking at the realized 3-month return.

To avoid threshold effects, we transform the final prediction for each stock into a score, where $+1$ stands for significantly positive excess return prediction, -1 for negative prediction and 0 for an uncertain prediction. The case where $S = 0$ is usually related to stocks for which some of their ESG indicators would eventually signal financial outperformance, while other ESG indicators rather signal potential underperformance. The picture is then nuanced, and the algorithm cannot make a precise prediction. This is a very common situation in finance, where different indicators can yield different forecasts, so that, in aggregate, the forecast turns out to be uninformative.

The learning process is divided into two steps. Following Nemirovski (2000) and Tsybakov (2003), the training set D_n at the end of year N is divided into two sub-datasets: D_n the learning set and D_t the aggregation set, with $t \gg n$ and $n + t = N$. The learning set D_n is used to design and select the set of rules used by the algorithm to make predictions. The aggregation set is used to fit the coefficients of the convex combination, in line with the expert aggregation theory of Cesa-Bianchi and Lugosi (2006) and Stoltz (2010).

4.1 Independent Suitable Rules

Let $D_N = ((x_1, y_1), \dots, (x_N, y_N)) \in \mathbb{X} \times \mathbb{R}^N$ be the training set. Here y_i denotes the 3-month return for some stock and x_i is the d -dimensional vector of its ESG features. The training set consists of a large but finite numbers of $(d + 1)$ -vectors spanning all stocks in the investment universe and all available dates. The training set $D_n \subseteq D_N$ includes the first n data points in D_N and $D_t = ((x_{n+1}, y_{n+1}), \dots, (x_N, y_N))$, the order being induced by the time.

For any set $E \subseteq \mathbb{X}$ we define:

$$\mu(E, D_n) := \frac{\sum_{i=1}^n y_i 1_{x_i \in E}}{\sum_{i=1}^n 1_{x_i \in E}} \quad (1)$$

Here, by convention, $\frac{0}{0} = 0$.

⁴ For each raw indicator, as for example the environment score (E), we also look at the derived indicator relative to the peer-group and the sector. All of these transformations can potentially contain useful information. On the other side, the use of both raw and derived indicators rapidly increases the dimension of the feature space d .

⁵ We use 164 ESG raw indicators, from which we derive peer group and sector relative indicators and 3 valuation indicators. In total $164 \times 3 + 3 = 495$. From these indicators we remove 48 indicators for which either the sector or the peer group derived indicators are too close, or for which historical data is missing.

The set-valued map μ represents the conditional excess return of a stock over the benchmark, given that its ESG features x belong to E .

Let r be a hyper-rectangle on \mathbb{X} : $r = \prod_{k=1}^d I_k$ where each I_k is an interval of X_k . A rule f is a function defined on $r \times (\times \mathbb{R})^n$ as:

$$f(x, D_n) := \mu(r, D_n) \forall x \in r \tag{2}$$

The hyper-rectangle r is called the condition and $\mu(r, D_n)$ is called the prediction of the rule f . The event $\{x \in r\}$ is called the activation conditions of the rule f .

A rule f is completely defined by its condition r . So, with an abuse of notation, we do not distinguish between a rule and its condition. We define two crucial numbers for a rule:

Let f be a rule defined on $r = \prod_{k=1}^d I_k$.

1. The number of activations of f in the sample D_n is

$$n(r, D_n) := \sum_{j=1}^n 1_{x_j \in r}$$

2. The complexity of f is

$$cp(r) = d - \#\{1 \leq k \leq d; I_k = X_k\}$$

The algorithm does not consider all the possible rules, but only those with a given coverage and significance. We call these rules suitable, and their definition is given below.

A rule f , defined on r , is a suitable rule for the training set D_n if and only if it satisfies the two following conditions:

1. Coverage condition.

$$C_{min} \leq \frac{n(r, D_n)}{n} \leq C_{max} \tag{3}$$

with C_{min} and C_{max} are suitably chosen in the calibration step.

2. Significance condition.

$$|\mu(r, D_n) - \mu(\mathbb{X}, D_n)| \geq z(r, D_n, \alpha) \tag{4}$$

For a chosen $\alpha \in [0, 1]$ and function z .

The coverage condition (3) excludes rules that are activated only on small sets (i.e. with a low coverage rate, C_{min}) and rules that are too obvious (i.e. with a high coverage rate, C_{max}). The threshold in the significance condition (4) is set such that the probability of falsely rejecting the null hypothesis $\mu(r, D_n) = \mu(\mathbb{X}, D_n)$ is less than α . The parameter α permits to control the number of suitable rules. The higher α , the higher the number of suitable rules. In what follows, we generate rules of complexity $c \geq 2$ by a suitable intersection of rules of complexity 1 and rule of complexity $c - 1$.

Two rules f_i and f_j defined on r_i and r_j respectively, form a suitable intersection if and only if they satisfy the two following conditions:

1. Intersection condition:

$$\begin{aligned} & r_i \cap r_j \neq \emptyset, \\ & n(r_i \cap r_j, D_n) \neq n(r_i, D_n) \neq n(r_j, D_n) \end{aligned} \tag{5}$$

2. Complexity condition:

$$cp(r_i \cap r_j) = cp(r_i) + cp(r_j) \tag{6}$$

The intersection condition (5) avoids adding a useless condition for a rule. In other words, to define a suitable intersection, r_i and r_j must not be satisfied by the same points in D_n . The complexity condition (6) means that r_i and r_j have no marginal index in common.

4.2 Designing Suitable Rules

The design of suitable rules is made recursively on their complexity. It stops at a complexity c if no rule is suitable or if the maximal complexity $cp = cp_{max}$ is achieved.

Complexity 1: The first step is to find suitable rules of complexity 1. First notice that the complexity of evaluating all rules of complexity 1 is $o(ndm^2)$. Rules of complexity 1 are the base of the algorithm search heuristic. So, all rules are considered and only suitable ones are kept, i.e. rules that satisfied the coverage condition (3) and the significance condition (4). Since rules are considered independently, the search can be parallelized.

Complexity c : Among the suitable rules of complexity 1 and $c - 1$, we select M rules of each complexity (1 and $c - 1$) according to a chosen criterion. Then it generates rules of complexity c by pairwise suitable intersection. The complexity of evaluating all rules of complexity c , obtained from their intersections, is $o(nM^2)$. Here again, since rules are considered independently, the evaluation can be parallelized. The parameter M helps to control the computing time.

4.3 Selecting Suitable Rules

We select a subset S from all suitable rules which maximizes the gains expected from rule in D_n and such as their conditions form a covering of \mathbb{X} .

4.4 Algorithm

The calibration of the algorithm is structured in two parts: in the first one, it finds all suitable rules, and in the second one it retains only an optimal subset of it. To avoid threshold effects, overfitting and to manage the numerical complexity, we discretize each feature in $\text{int}(\mathbb{X}/m)$ classes with empirical quantiles (modalities)⁶. Thus, each modality of each variable covers about $100/m$ percent of the sample. In practice, m must be inversely related to d : The higher the dimension of the problem, the smaller the number of modalities.

The parameters of the algorithm are:

- m , the sharpness of the discretization;
- $\alpha \in [0, 1]$, which specifies the false rejecting rate of the test;
- z , the significance function of the test;
- C_{max} and C_{min} the coverage bounds;
- cp_{max} the maximal complexity of a rule;
- and $M \in \mathbb{N}$, the number of rules of complexity 1 and $c - 1$ used to define the rules of complexity c .

4.5 Aggregation

Let $D_t = ((x_{n+1}, y_{n+1}), \dots, (x_N, y_N)) \in \mathbb{X} \times \mathbb{R}^t$ be the aggregation set and let S be the set of R rules selected by the algorithm. At each time t , the predictions of each rule are aggregated into one prediction \hat{y}_t as follows:

$$\hat{y}_t = \frac{\sum_{i=1}^R \pi_{i,t} f_i(x_t, D_n)}{\sum_{i=1}^R \pi_{i,t} 1_{x_t \in r_i}} \quad (7)$$

where $\pi_{i,1} = 1/R$. When the realized value y_t is known, the weights $\pi_{i,t+1}$ are updated with the following formula:

$$\pi_{i,t+1} = \pi_{i,t} \frac{e^{-\eta \ell(f_i(x_t, D_n), y_t)}}{\sum_{i=1}^R e^{-\eta \ell(f_i(x_t, D_n), y_t)}}, \quad (8)$$

where $\eta > 0$ and ℓ convex loss function.

Once trained, the machine learning algorithm produces predictions of the excess returns which are transformed into a scores $S \in \{-1, 0, +1\}$, given the out-of-sample ESG features x_t for each company. Table 9 shows some examples of rules taken from the learning process of the algorithm. The table lists three rules associated with positive predictions (opportunities) and five rules with negative predictions. Each rule consists of two features and two intervals. The "Relative To" properties indicate whether the feature must be calculated over all stocks in the universe (All), over a Sector, over a Peer Group, or whether we should look at the variations of the feature over time (Delta Score).

⁶ Of course, such procedure is performed only on real-valued features with more than m different values. Categorical features are left unchanged.

Whenever the values taken by the features for a given company fall in the given intervals (we say that the stock activates the rule) the algorithm makes a prediction on its excess return. It is important to remark that we aggregate all the predictions, and we transform the final aggregated prediction into a score $\{-1, 0, +1\}$, so that in the end we mainly look at the sign of the prediction rather than at its magnitude. We also remark that, while the set of rules remains unchanged for one year (until the next learning process), the output of the rules can change over time, because raw indicators can change and also because the aggregated weights of the rules change over time.

5. Machine Learning Application

We now test the predictive power of the machine learning algorithm developed in Section 4 compared with the classical best-in-class approach. More precisely, we try to assess whether filtering stocks over scores derived from the algorithm outperforms the standard filtering over ESG ratings (best-in-class). For the sake of simplicity, we only present the World Developed universe and, among the strategies presented in Section 3, we only consider the 30% best-in-class, as it is very close to what investors look at for their ESG portfolios. We recall that this strategy excludes, at each monthly review, the stocks whose ESG ratings are in the lower tercile within each peer group, and finally scale the weights so that their sum is one. To ensure replicability of the portfolio, the ESG ratings are taken four days before the review date (which is end-of-month).

At the monthly review, we also build three portfolios based on the scores calculated with the machine learning algorithm, with the rules calculated at the end of the year that precedes the review:

Positive ML Screening: The portfolio selects all stocks in the investment universe whose scores are $+1$. The weights are finally scaled up to sum to one (maintaining then the capitalization-weighting scheme of the benchmark)

Positive ML Screening Sector Matched: Same selection as for the Positive ML Screening portfolio, but the scaling of the weights is done in such a way that the final sector breakdown of the portfolio is matched to the benchmark's one.

Negative ML Screening: The portfolio selects all stocks in the investment universe whose scores are -1 . The weights are finally scaled up to sum to one (maintaining then the capitalization-weighting scheme of the benchmark)

Table 9. Examples of rule

Feature	Relative	Activation	Rule Description
	To	Set	
Business Ethics Incidents	Sector	[5, 9]	WHEN Business Ethics Incidents is high relative to sector AND Board Remuneration Disclosure is high relative to sector THEN Opportunity
Board Remuneration Disclosure	Sector	[5, 9]	
Board Independence	All	[9, 9]	WHEN Board Independence is at the maximum AND Board Remuneration Disclosure is high relative to sector THEN Opportunity
Board Remuneration Disclosure	Sector	[5, 9]	
Max. Verification of ESG Reporting	Sector	[0, 7]	WHEN Verification of ESG Reporting is not high relative to sector AND Board Remuneration Disclosure is low relative to sector THEN Risk
Board Remuneration Disclosure	Sector	[0, 4]	
Quantitative Performance	All	[5, 9]	WHEN Quantitative Performance Score is high AND Board Remuneration Disclosure is low relative to sector THEN Risk
Board Remuneration Disclosure	Sector	[0, 4]	

Description: Some rules from the learning process of the algorithm at the end on 2012, 2013 and 2016. All features are discretized over 10 modalities (0 to 9) except for Quantitative Performance which is discretized over 6 modalities (0 to 5). High values for the features correspond to good ESG performance.

Table 10. Key performance indicators

	ML Screening			ESG best_in_class	
	Bench	Positive	Positive Sect. Matched	Negative	30%
Ann. Performance	10.32%	13.07%	11.66%	8.31%	10.13%
Ann. Volatility	10.50%	11.14%	10.96%	10.95%	10.57%
Sharpe Ratio	0.94	1.14	1.03	0.72	0.92
Max. Drawdown	-18.07%	-14.99%	-16.46%	-22.47%	-17.91%
Information Ratio	-	1.01	0.58	-0.54	-0.32
Ann. Alpha	-	2.47%	1.15%	-1.81%	-0.24%

Description: Key performance indicators of the MSCI World Index (Bench.), the capitalization-weighted selection filtered over positive scores from the ML algorithm, the one with the sector allocation matched to the benchmark, the one screened over negative scores and the 30% ESG best-in-class filtered portfolios. Data is shown in USD from January 2013 to March 2018. Source MSCI, Datastream, Sustainalytics.

As before, the scores are taken four days before the review date. We consider the sector matched portfolio because the absolute screening usually introduces significant sector deviations with respect to the benchmark. Table 10 collects the main results for these portfolios since January 2013.

Although we recognize that the period over which we can test the machine learning algorithm is relatively short (five years and three months), the results we obtain contain some interesting insights. First of all, the Positive ML Screening outperforms all the other portfolios: by 2.76% the benchmark on an annualized basis, by 2.94% the ESG best-in-class portfolio and by 4.77% the Negative ML Screening. And while the realized annual volatilities remain in the range 10.50% to 11.14%, there are significant differences in the realized maximum drawdowns: the Negative ML Screening shows a -22.47% loss from its peak, while the Positive ML Screening loss from its peak accounts for -14.99% .

These two combined results show that the machine learning algorithm is clearly able to distinguish between opportunity stocks (the ones with positive scores) from risky stocks (negative scores). Figure 1 shows the historical behaviour of these two portfolios and the benchmark. We notice that the Positive ML Screening outperforms the Negative ML Screening over time, with the benchmark in between. Furthermore, in years when the benchmark shows very high performances with very low volatility, typically in bull market regimes, the differences between the two strategies are less pronounced. On the contrary, when the market is in bear regimes or it does not have clear trend, the Positive ML Screening clearly outperforms its negative counterpart, as shown in Table 11.

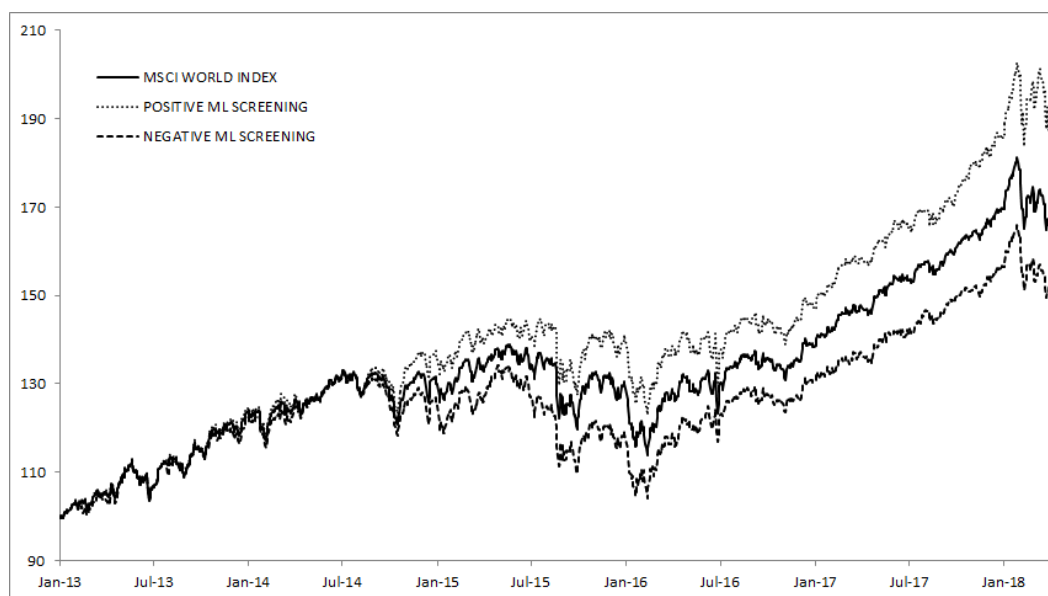


Figure 1. Simulated strategy

Description: Simulated strategy levels for the benchmark MSCI World Index, the capitalization-weighted selection filtered over positive scores from the ML algorithm (Positive ML Screening) and the one screened over negative scores (Negative ML Screening). Data in USD from January 2013 to March 2018. Base level = 100. Source MSCI, Datastream, Sustainalytics.

In years when the benchmark performance is very significant (2013 or 2017), the Positive ML Screening is still able to achieve some outperformance, but the spread with the Negative ML Screening is somehow lower than years when the market performance is negative or low (2014, 2015 and, most recently, 2018).

Interestingly, the excess return of the sector matched version is also positive, even if lower in magnitude when compared to the Positive ML Screening. By neutralizing the sector component (because matched), the outperformance essentially comes from the stock picking.

Table 11. Calendar year performances

Year	Bench	Excess Return			ESG best_in_class 30%
		Positive	Positive Sect. Matched	Negative	
2013	23.95%	0.72%	0.12%	-1.39%	-0.14%
2014	4.97%	3.88%	3.65%	-2.75%	-0.17%
2015	-0.89%	3.79%	1.89%	-5.3%	0.07%
2016	7.4%	-2.14%	-1.91%	3.73%	-0.02%
2017	22.44%	3.82%	0.61%	-2.57%	-0.02%
2018	-1.37%	3.92%	2.24%	-1.63%	-0.24%

Description: Calendar year performances for the MSCI World Index (Bench.) and the excess returns for the Positive, Positive Sect. Matched and Negative ML Screening as well as for the ESG 30% best-in-class portfolio. Data is shown in USD from January 2013 to March 2018. Source MSCI, Datastream, Sustainalytics.

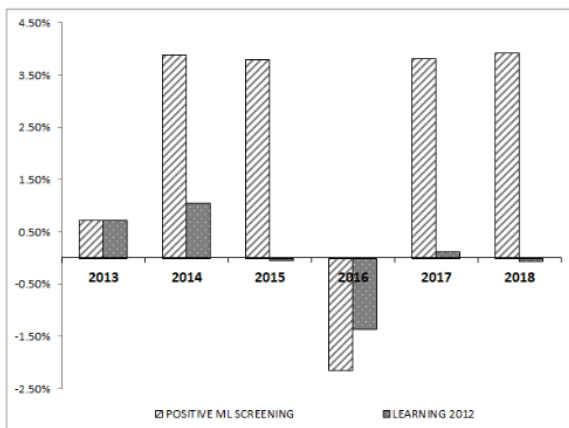
For the Negative ML Screening, the excess return is always negative except for 2016. Finally, the best-in-class portfolio shows almost systematically small but negative excess returns, except in 2015 when it managed to outperform by 0.07%. Once again, our findings confirm the fact that for very large and diversified universes, the simple ESG filtering does not bring alpha, although it does not significantly reduce the performance with the best-in-class approach.

5.1 The Effects of Learning

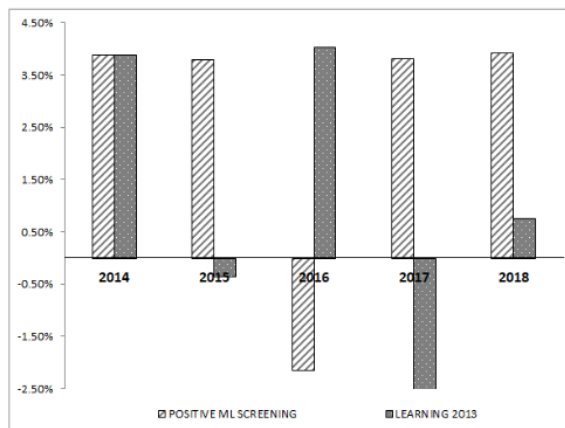
The machine learning algorithm is initially trained over three years of data and then yearly updated. During these regular updates, the algorithm learns from the new flow of data it can access: It can test its rules to confirm, nuance or remove some of them, and selects new rules linked to statistically significant patterns. This learning process is key in the final performance of the model (and for the Positive ML Screening portfolio built upon it). To measure this effect, we form four portfolios named LEARNING Y , where $Y = 2012, 2013, 2014, 2015$ as follows:

- For each year Y , we consider the set of rules related to the learning at the end of the year Y .
- We calculate the scores for all stocks in the universe from the end of year Y to March 2018 with this set of rules.
- LEARNING Y is built as Positive ML Screening, except that the underlying scores are calculated with the same, not updated set of rules calibrated at the end of year Y .

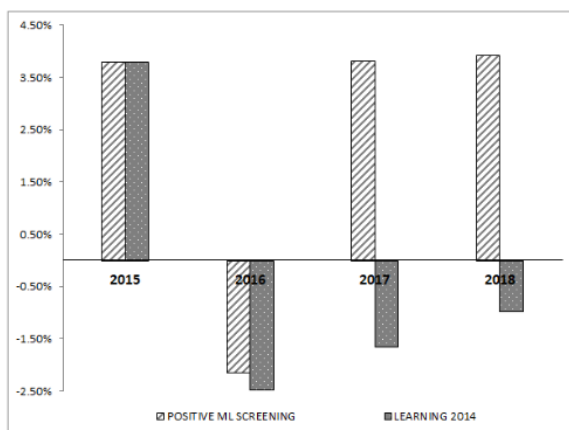
Said differently, LEARNING Y uses a unique, static set of rules that is never updated (no learning). By construction, the portfolios Positive ML Screening and LEARNING Y coincide over the period January, 1st, $Y + 1$ to December, 31st, $Y + 1$, because, over this period, they use the same set of rules (hence the same scores) to screen the investment universe. Figure 2 shows the calendar excess returns of these portfolios together with the Positive ML Screening portfolio over the benchmark MSCI World Index.



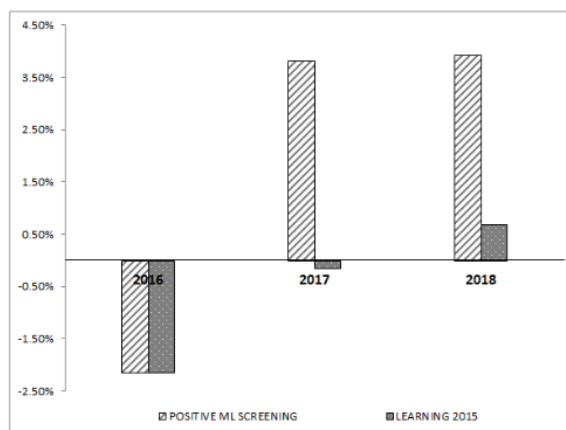
(a) Learning end of 2012



(b) Learning end of 2013



(c) Learning end of 2014



(d) Learning end of 2015

Figure 2. Calendar excess returns

Description: Calendar excess returns of the Positive ML Screening and the four portfolios LEARNING 2012, LEARNING 2013, LEARNING 2014 and LEARNING 2015 over the MSCI World Index. Data is shown in USD from January 2013 to March 2018. Source MSCI, Datastream, Sustainalytics.

Since we only show out-of-sample results, the time frame of each LEARNING Y portfolio is different. In the majority of cases, we see that Positive ML Screening outperforms the LEARNING Y portfolios after the first year (since they are the same on the first year).

Indeed, the excess return for the LEARNING Y portfolios usually shrinks to zero and becomes even negative over time.

In other words, the predictive power of the scores vanishes over time, so that it is important to train the algorithm on the new observed data to update the set of rules.

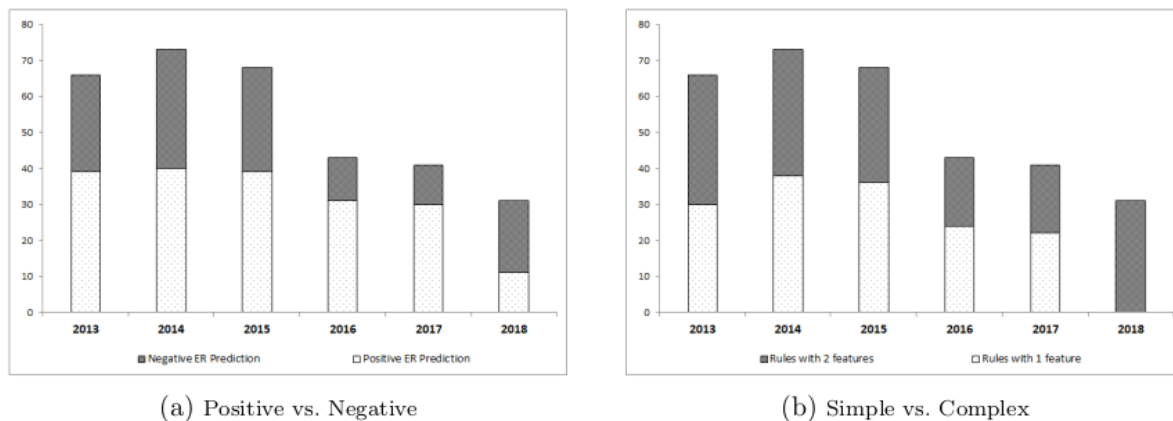


Figure 3. Number of rules at each update of the algorithm

Description: Number of rules at each update of the algorithm: (a) the split between rules that predict positive or negative excess returns (ER); (b) the split between rules that make use of one feature (Simple) or two features (Complex).

The number of rules used by the algorithm changes over time: as shown in Panel (a) of Figure 3, this number evolves in the range [31, 73] with the split between positive rules (i.e. rules related to positive predictions of the excess return) and negative ones also changing over time. Interestingly, the number of rules related to negative excess return increased from 12 in 2016 to 20 in the latest 2018 learning. Panel (b) of Figure 3 shows the same number of rules split between simple rules (i.e. those that only make use of one feature) and complex rules (i.e. those that use two features, as the examples shown in Table 9). Both Figures 2-3 suggest that to extract alpha from the ESG features, one needs to regularly update the algorithm, and consider newly created set of rules to detect patterns between ESG profiles and financial performances.

6. Conclusion

The last few years have seen an increasing interest toward ESG investing and the integration of socially responsible principles at the portfolio construction level. Managers and investors are asked to complement pure financial objectives with extra-financial ones.

Our study brings some new ideas and insights onto the way investors could achieve ESG objectives in their investments. The literature on the theme is mixed: Initial studies were mostly skeptical on the benefit of ESG integration into the portfolio. Over time the mindset has evolved, and several studies have empirically proved that ESG integration in the portfolio does not lower performances. Most recently, the financial literature has gone one step further and claim that, indeed, ESG integration is a way to extract alpha or, at least, to reduce risks.

We do recognize the need for serious integration of ESG objectives alongside with classic financial ones, and that there exists an economic link between the ESG profile of a company and its financial performances over the long run. Nevertheless, we tend to agree with the pioneers of ESG research for which, at best, ESG integration does not significantly degrade financial performances, especially for large and diversified investment universes.

Because ESG profiles can impact financial performances in a non-linear way, and the impact can depend on the sector, the country or other specific characteristics of each company, we designed and implemented a sophisticated machine learning algorithm that identifies patterns between ESG profiles and performances, statistically robust across the universe and over time.

The algorithm produces a set of rules, each rule identifying a region in the high-dimensional space of the ESG features, conditionally on which we can make a prediction on the stock's excess return. All the predictions are finally aggregated and transformed into a score taking values in $\{-1, 0, +1\}$, so that in the end we effectively look at the sign of the excess return rather than its magnitude.

With this algorithm, trained over time to keep it updated, we empirically proved that the link between ESG profiles and financial performances exists, but can only be accessed with non-linear techniques. Indeed, a simple strategy that selects stocks whose scores are positive significantly outperforms the ESG best-in-class approach.

Annexes

Table 1. World Developed

	ESG best-in-class			
	Bench.	10%	30%	50%
Ann. Performance	10.07%	10.01%	9.93%	9.51%
Ann. Volatility	13.34%	13.31%	13.44%	13.82%
Sharpe Ratio	0.73	0.73	0.72	0.67
Max. Drawdown	-21.91%	-21.79%	-22.02%	-22.57%
Information Ratio	0	-0.27	-0.25	-0.41

Table 2. US

	ESG best-in-class			
	Bench.	10%	30%	50%
Ann. Performance	13.45%	13.25%	13.46%	13.2%
Ann. Volatility	14.61%	14.54%	14.49%	14.4%
Sharpe Ratio	0.9	0.89	0.91	0.9
Max. Drawdown	-18.99%	-18.87%	-18.71%	-18.04%
Information Ratio	0	-0.71	0.02	-0.18

Table 3. Europe

	ESG best-in-class			
	Bench.	10%	30%	50%
Ann. Performance	6.37%	6.55%	6.47%	6.31%
Ann. Volatility	19.25%	19.19%	19.19%	19.29%
Sharpe Ratio	0.32	0.33	0.32	0.31
Max. Drawdown	-30.25%	-30.21%	-30.2%	-30.54%
Information Ratio	0	0.43	0.22	-0.11

Table 4. Asia

	ESG best-in-class			
	Bench.	10%	30%	50%
Ann. Performance	6.83%	6.71%	6.41%	5.75%
Ann. Volatility	15.54%	15.71%	16%	16.2%
Sharpe Ratio	0.42	0.41	0.38	0.34
Max. Drawdown	-24.8%	-24.95%	-25.27%	-25.8%
Information Ratio	0	-0.21	-0.36	-0.46

Description: Key performance indicators of the MSCI World Index and three capitalization-weighted regional benchmarks, together with ESG best-in-class filtered portfolios with different thresholds: 10%, 30% and 50%. Data is shown in USD from August 2009 to March 2018. Source MSCI, Datastream, Sustainalytics.

Table 5. World Developed

	ESG	E	S	G
CS	-0.26%	-0.45%	-0.65%	0.62%
	(-0.26)	(-0.47)	(-0.59)	(0.53)
CD	-0.33%	-0.46%	-0.65%	-1.22%
	(-0.3)	(-0.46)	(-0.4)	(-0.65)
EN	-0.24%	-0.41%	-0.26%	0.86%
	(-0.13)	(-0.19)	(-0.17)	(0.31)
FI	-0.51%	-0.46%	-0.87%	0.19%
	(-0.43)	(-0.42)	(-0.63)	(0.13)
HC	-0.39%	-0.5%	-0.47%	-0.53%
	(-0.35)	(-0.3)	(-0.4)	(-0.43)
IN	-0.31%	-0.32%	-0.32%	-0.33%
	(-0.28)	(-0.27)	(-0.31)	(-0.19)
IT	-0.13%	-0.44%	-1.53%	0.1%
	(-0.12)	(-0.45)	(-0.52)	(0.05)
MA	-0.09%	-0.09%	-0.17%	0.32%
	(-0.05)	(-0.03)	(-0.08)	(0.18)
TL	1.26%	1.17%	0.15%	0.73%
	(0.64)	(0.74)	(0.06)	(0.21)
UT	0.16%	-0.94%	0.72%	0.82%
	(0.1)	(-0.39)	(0.43)	(0.32)

Table 6. US

	ESG	E	S	G
CS	-0.94%	-0.46%	-0.46%	-0.26%
	(-0.84)	(-0.43)	(-0.32)	(-0.12)
CD	-1.95%	-0.48%	-0.91%	-2.92%
	(-0.82)	(-0.39)	(-0.55)	(-0.98)
EN	-0.09%	-0.66%	0.01%	0.55%
	(-0.06)	(-0.44)	(0.01)	(0.16)
FI	-0.22%	-0.11%	-0.2%	1.02%
	(-0.19)	(-0.11)	(-0.17)	(0.67)
HC	-0.23%	0.19%	-0.55%	-0.58%
	(-0.22)	(0.15)	(-0.5)	(-0.37)
IN	0.4%	0.25%	0.01%	-0.6%
	(0.54)	(0.32)	(0.01)	(-0.56)
IT	-1.36%	-0.7%	-1.76%	-0.37%
	(-0.96)	(-0.75)	(-0.51)	(-0.11)
MA	0.34%	0.72%	-0.76%	-0.14%
	(0.18)	(0.32)	(-0.32)	(-0.07)
TL	-0.38%	-0.32%	-0.4%	1.32%
	(-0.37)	(-0.38)	(-0.2)	(0.28)
UT	0.63%	0.47%	0.7%	0.13%
	(0.56)	(0.38)	(0.7)	(0.12)

Table 7. Europe

	ESG	E	S	G
CS	0.16% (0.14)	-0.03% (0.03)	-0.01% (-0.01)	0.39% (0.32)
CD	0.51% (0.32)	0.45% (0.29)	0.18% (0.14)	1.07% (0.68)
EN	-1.2% (-0.3)	-2.01% (-0.38)	-0.5% (-0.12)	-1.05% (-0.27)
FI	0.31% (0.18)	0.04% (0.03)	0.33% (0.17)	-0.28% (-0.21)
HC	-0.38% (-0.49)	-0.44% (-0.55)	-0.53% (-0.67)	-0.11% (-0.18)
IN	0.03% (0.03)	0.03% (0.03)	-0.39% (-0.35)	-0.67% (-0.44)
IT	-0.63% (-0.29)	-1.13% (-0.49)	0.01% (0.08)	-0.71% (-0.43)
MA	0.12% (0.03)	0.23% (0.06)	0.51% (0.14)	0.6% (0.18)
TL	1.81% (0.67)	1.72% (0.63)	1.82% (0.62)	2.06% (0.71)
UT	1.79% (0.54)	3.25% (0.87)	0.09% (0.03)	0.41% (0.16)

Table 8. Asia

	ESG	E	S	G
CS	0.6% (0.52)	0.37% (0.12)	-0.3% (-0.28)	-0.18% (-0.09)
CD	-0.29% (-0.19)	-0.11% (-0.11)	0.09% (0.04)	0.75% (0.37)
EN	0.12% (0.04)	0.34% (0.11)	0.03% (0.01)	0.25% (0.1)
FI	-0.27% (-0.16)	-0.09% (-0.06)	-0.71% (-0.46)	-0.08% (-0.05)
HC	0.29% (0.19)	-0.67% (-0.27)	0.14% (0.08)	-0.03% (-0.02)
IN	-0.76% (-0.39)	-0.41% (-0.22)	-0.42% (-0.22)	-0.61% (-0.35)
IT	-1.32% (-0.59)	-0.92% (-0.4)	-1.01% (-0.43)	-0.96% (-0.28)
MA	-0.09% (-0.05)	-0.09% (-0.03)	-0.17% (-0.08)	0.32% (0.18)
TL	-0.29% (-0.26)	-0.18% (-0.11)	-0.61% (-0.09)	1.83% (0.24)
UT	3.07% (0.65)	-3.57% (-0.87)	0.16% (0.04)	2.6% (0.55)

Description: Annualized excess returns (Information ratios) between capitalization-weighted sector portfolios and their ESG best-in-class filtered versions for the MSCI World Index and the derived regional benchmarks. In bold pairs sector/indicator for which the excess return is positive. Best-in-class filters are performed with the ESG rating together with the single pillars Environment (E), Social (S) and Governance (G) ratings. Data is shown in USD from August 2009 to March 2018. Source MSCI, Datastream, Sustainalytics.

References

- Allouche, J., & Laroche, P. (2005). A meta-analytic investigation of the relationship between corporate social and financial performance. *Revue de gestion des ressources humaines*, 57, 18.
- Asness, C. (2017). Virtue is its Own Reward: Or, One Mans Ceiling is Another Mans Floor. AQR Blog. Retrieved from <https://www.aqr.com/Insights/Perspectives/Virtue-is-its-Own-Reward-Or-One-Mans-Ceiling-is-Another-Mans-Floor>
- Aupperle, K. E., Carroll, A. B., & Hatfield, J. D. (1985). An empirical examination of the relationship between corporate social responsibility and profitability. *Academy of management Journal*, 28(2), 446-463.
- Bragdon, J. H., & Marlin, J. A. T. (1972). Is pollution profitable? *Risk Management*, 19(4), 9-18.

- Capelle-Blancard, G., & Monjon, S. (2012). Trends in the literature on socially responsible investment: Looking for the keys under the lamppost. *Business Ethics: A European Review*, 21(3), 239-250.
- Cesa-Bianchi, N., & Lugosi, G. (2006). *Prediction, Learning and Games*. Cambridge university press.
- Chong, J., & Phillips, G. M. (2016). ESG investing: A simple approach. *The Journal of Wealth Management Fall 2016*, 19(2), 73-88.
- Filbeck, G., Holzhauser, H. M., & Zhao, X. (2014). Using social responsibility ratings to outperform the market: Evidence from long-only and active-extension investment strategies. *The Journal of Investing Spring 2014*, 23(1), 79-96
- Friede, G., Bush, T., & Bassen, A. (2015). ESG and financial performance: aggregated evidence from more than 2000 empirical studies. *Journal of Sustainable Finance & Investment*, 5(4), 210-233
- Friedman, M. (1970). *The social responsibility of business is to increase its profits*. The New York Times Magazine, September 13, 1970. Copyright 1970 by The New York Times Company.
- Giese, G., Ossen, A., & Bacon, S. (2016). ESG as a performance factor for smart beta indexes. *The Journal of Index Investing Winter 2016*, 7(3), 7-20.
- Griffin, J. J., & Mahon, J. F. (1997). The corporate social performance and corporate financial performance debate: Twenty-five years of incomparable research. *Business & society*, 36(1), 5-31.
- Humphrey, J. E., & Tan, D. T. (2014). Does it really hurt to be responsible? *Journal of business ethics*, 122(3), 375-386.
- Indrani De, I., & Clayman, M. R. (2015). The benefits of socially responsible investing: An active manager's perspective. *The Journal of Investing Winter 2015*, 24(4), 49-72.
- Kurtz, L., & Di Bartolomeo, D. (2011). The long-term performance of a social investment universe. *The Journal of Investing Fall 2011*, 20(3), 95-102.
- Margolis, J. D., Elfenbein, H. A., & Walsh, J. P. (2009). *Does it pay to be good and does it matter? A meta-analysis of the relationship between corporate social and financial performance*.
- Peiris, D., & Evans, J. (2010). The relationship between environmental social governance factors and U.S. stock performance. *The Journal of Investing Fall 2010*, 19(3), 104-112.
- Nemirovski, A. (2000). Topics in Nonparametric. *Ecole d'Été de Probabilités de Saint-Flour*, 28, 85.
- Orlitzky, M., Schmidt, F. L., & Rynes, S. L. (2003). Corporate social and financial performance: A meta-analysis. *Organization Studies*, 24(3), 403-441.
- Revelli, C., & Viviani, J. L. (2015). Financial performance of socially responsible investing (SRI): What have we learned? A meta-analysis. *Business Ethics: A European Review*, 24(2), 158-185.
- Shiller, R. J. (2013). Capitalism and financial innovation. *Financial Analysts Journal*, 69(1).
- Stoltz, G. (2010). Agrégation séquentielle de prédicteurs: méthodologie générale et applications à la prévision de la qualité de l'air et celle de la consommation électrique. *Journal de la Société Française de Statistique*, 151(2), 66-106.
- Tsybakov, A. B. (2003). *Optimal Rates of Aggregation*. Learning Theory and Kernel Machines, Springer, pp. 303-313.
- Van Beurden, P., & Gössling, T. (2008). The worth of values a literature review on the relation between corporate social and financial performance. *Journal of Business Ethics*, 82(2), 407-424.
- Van Duuren, E., Plantinga, A., & Scholtens, B. (2016). ESG Integration and the Investment Management Process: Fundamental Investing Reinvented. *Journal of Business Ethics*, 138, N. 3, pp. 525 - 533
- Wu, M.L. (2006). Corporate social performance, corporate financial performance, and firm size: A meta-analysis. *Journal of American Academy of Business*, Vol. 8, N. 1, pp. 163 - 171
- Zoltan, N. and Kassam, A. Lee, L.E. (2016). Can ESG add alpha? An analysis of ESG tilt and momentum strategies. *The Journal of Investing Summer 2016*, Vol. 25, N. 2, pp. 113 - 124

Copyrights

Copyright for this article is retained by the author(s), with first publication rights granted to the journal.

This is an open-access article distributed under the terms and conditions of the [Creative Commons Attribution license](#) which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.