
ASSESSING RISK THROUGH ENVIRONMENTAL, SOCIAL AND GOVERNANCE EXPOSURES

Jeff Dunn^{a,b}, Shaun Fitzgibbons^{a,c} and Lukasz Pomorski^{a,d}

We discuss risk and return implications of incorporating environmental, social and governance (ESG) considerations in an investment strategy and argue that ESG exposures may be informative about the risks of individual firms. We show empirically that stocks with worst ESG exposures have volatility that is up to 15% higher, and betas up to 3% higher, than stocks with the best ESG exposures. We also find that ESG scores may help forecast future changes to risk estimates from a traditional risk model. Controlling for the contemporaneous risk model estimates, we show that poor ESG exposures predict increased future statistical risks.



Asset owners and allocators are increasingly interested in the Environment, Social and Governance (ESG) profile of their investments. While some investors have always incorporated ethical or sustainable themes into their portfolios, the amount of assets managed in ESG-conscious mandates has increased dramatically in recent years. For example, the US Social Investment Forum reports in “US Sustainable, Responsible

and Impact Investing Trends, 2014” that the total US-domiciled assets grew from \$3.74 trillion in 2012 to \$6.57 trillion in 2014 and that the 2014 level of assets represents a sixth of the total assets under professional management in the US. One obvious driver of this trend is ethical in nature: investors have concerns for the environment, unease about investing in “sin stocks,” etc. We acknowledge such motives, but do not discuss them in this paper. Instead, we focus on the portfolio implications of ESG-informed investing. Such implications may affect either risk or returns. We begin with a general overview of both, and argue that it is more plausible that ESG correlates with a dimension of risk rather than that it influences expected returns.¹ Afterward, we focus exclusively on risk in our paper.

^aAQR Capital Management, LLC, 2 Greenwich Plaza, Greenwich, CT, 06830.

^bPhone: (+61 2) 8023-6501. E-mail: Jeff.Dunn@aqr.com

^cPhone: (+1 203) 742-3844. E-mail: Shaun.Fitzgibbons@aqr.com

^dPhone: (+1 203) 742-6687. E-mail: Lukasz.Pomorski@aqr.com

Our main contribution is an empirical investigation of the potential link between ESG exposures and risk. In an effort to be transparent and allow others to easily replicate our results we utilize well-known, third-party measures of ESG and risk: MSCI ESG database for the former and Barra's risk models for the latter. We find strong evidence that the two are interrelated and that stocks with poor ESG profiles are riskier according to the statistical risk model. Stocks in the worst ESG quintile have total volatility and stock-specific volatility that is higher by 10–15%, and betas that are higher by 3%, than the corresponding measures for stocks in the best ESG quintile. This pattern is not only robust to a variety of controls, but is also clear over time and within various investment universes, for example in US equities, in developed markets outside of the US, and in emerging markets.

Finding that ESG measures correlate with contemporaneous risk forecasts is an important piece of evidence. We go further and also document that poor ESG exposure tends to predict increases in statistical risks (i.e., risks captured by traditional risk models) in the future. Controlling for current risk characteristics of a given stock, that stock's ESG score helps forecast future statistical risks up to five years later. In other words, ESG exposures may convey information about future risks that are not captured by statistical risk models. The magnitude of the effect is modest: we find that deterioration of ESG score from the 75th to the 25th percentile is associated with about a 1% increase in risk. This may reflect the fact that ESG captures risks that are long-run in nature and may not materialize in short to medium term. For example, a firm with poor governance may be more likely to experience a scandal, earnings misstatement, etc., but that does not mean that such an event will necessarily happen over our relatively short sample period and consequently be captured in a statistical risk model. Moreover, we

are using a state-of-the-art risk model that already reflects much of the information about a given stock's risk, whether such risk is driven by ESG or any other type of exposure. Thus, it is by no means obvious that ESG information could help improve on this risk model's forecast, even if it is by a modest amount. Overall, we conclude that investors might be able to utilize ESG information to glean additional insights about the riskiness of their investments.

1 ESG characteristics and stocks' risks and returns

1.1 Risks

To date the literature on ESG has largely focused on the impact of ESG exposures on returns and little has been done to directly assess the impact on risk. If we consider risk as any form of uncertainty, and further recognize that ESG by its very nature is dealing with the impact of corporate activities on stakeholders, then it stands to reason that there is a direct link between management on the dimensions of E, S and G and the range of potential impacts on the stakeholders of a corporation. It is logical to postulate that companies neglecting to manage their ESG exposures may be exposed to higher risk (a wider range of potential outcomes) than their more ESG-focused counterparts.

To give a few simple examples, a firm that produces high levels of emissions during a manufacturing process may be exposed to potential future legislation that might impose a carbon tax; a firm poorly treating its employees or suppliers may face a backlash from its consumers and see its sales plummet; a firm with poor governance may get involved in a scandal that ultimately causes its downfall. Despite the diversity of situations that ESG issues can cover, these examples have a few characteristics in common: They each describe events that may have meaningful impact

on firm value, even though that impact is both uncertain in the timeframe over which it may be realized and potentially difficult to quantify or model. These concerns are of course not new, and have been recently raised by policymakers or by asset owners and allocators.²

The long term and infrequent nature of ESG events combined with the difficulty in modeling potential outcomes suggest that the risks reflected in ESG exposures may not be fully captured by traditional risk models, which are based on historical data and typically calibrated to shorter-term horizons. At the same time, if ESG captures some element of risk, we would expect it to show up at least partly in statistical risk models. If the potential risks were never realized and never picked up by risk models over long samples, then we may legitimately wonder whether there were any risks there to begin with. Thus, to the extent that ESG captures some dimension of risks, we would expect some correlation between ESG scores and traditional risk measures, perhaps not only contemporaneously but also in a predictive manner (ESG predicting what risk models may say in the future). We take these ideas to the data and test them in subsequent sections.

1.2 Returns

While our paper focuses on the relationship between ESG and risk, we also briefly review prior literature on the potential return implications of tilts towards better ESG stocks. Most academic studies argue that, if anything, stocks that rank poorly on ESG may bring relatively higher returns. The basic economic intuition is that some investors are unwilling to hold companies with poor ESG exposures, and that this reduction in the demand for shares may translate into lower prices today and higher returns in the future. In other words, investor preferences matter and investors' demand has an impact on stock prices, as discussed for example in Fama

and French (2007). This view is generally validated by empirical work. For example, Hong and Kacperczyk (2009) find that "sin stocks (. . .) have higher expected returns than otherwise comparable stocks." Similar evidence can be found, for example, in Fabozzi *et al.* (2008) or Statman and Glushkov (2009).

At the same time, another strand in the literature has suggested that certain dimensions of ESG, most notably governance, may correlate with better returns. For example, Gompers *et al.* (2003) construct a governance index (G-index) based on the number of provisions that may decrease shareholder rights (golden parachutes, staggered boards, etc.). The study documents that firms with higher values of the G-index and, presumably, poorer governance, realize lower average returns.

However, subsequent research has suggested that the relationship between governance and returns may not be quite as strong as previously believed (e.g., see the discussion in Larcker *et al.*, 2007). Similarly, Bebchuk *et al.* (2013) replicate the Gompers *et al.*'s (2003) study, finding similarly strong return patterns in the same sample (the 1990s) but no return predictability out of sample (the 2000s). Bebchuk *et al.* (2013) interpret as evidence of investors learning about stocks' ESG profile and incorporating that information into prices.

Given the somewhat mixed evidence on whether ESG could drive individual stock returns, it is perhaps not surprising that mutual fund studies tend to find similar performance for ESG-/SRI-focused funds and for conventional mutual funds. One recent example is Borgers *et al.* (2015). The study documents that mutual funds with higher sin stock exposure realize somewhat higher returns, but shows that the return differences tend not to be statistically significant.

2 Data

2.1 Measuring ESG exposures

To identify companies' ESG exposures we utilize the MSCI ESG database (often referred to as the Intangible Value Assessment, or IVA, data) over the period of January 2007 to December 2015.³ The IVA methodology covers a large cross-section of companies around the world (as of December 2015, over 5,000 companies, accounting for 97% of the market cap of the MSCI World index) and assesses how much each company is exposed to, and how well it manages its exposure to a range of environmental, social and governance issues affecting its industry. For each industry, IVA identifies key issues based on the extent to which businesses involved in that industry are exposed to or create large externalities in three areas:

- Environmental, including climate change, natural capital, pollution and waste;
- Social, including human capital, product liability, and stakeholder opposition; and
- Governance, including corporate governance and corporate behavior.

Within each of these areas, companies are assessed based on the combination of their individual business exposures, the magnitude of the risk, and the extent to which management has addressed the risk in their strategies and governance. As such, companies which have a large risk exposure but also have adequate measures in place to control the risks will not be ranked poorly. Moreover, companies that have exposure to only small risks will not rank poorly even if they have limited controls in place to manage such risks. Finally, IVA determines its final ratings on an industry-relative basis so that the resulting company scores can be more easily compared across various industries.

The IVA methodology aims to produce a holistic assessment of companies' ESG risk, including risk exposures based on their business segments, geographic segments, and risks specific to a particular company. To that end, IVA uses a range of data sources and documents, for example corporate filings, government data, news media, or relevant organizations and professionals.

2.2 Statistical measures of risk

Our investigation covers stocks in the US, international developed, and emerging markets. We select stocks that are members of the major indexes in these three equity universes: Russell 3000 for the US, MSCI World ex US for international developed, and MSCI Emerging for emerging markets.⁴

To measure the risks of these stocks, we utilize Barra's GEM2L risk model and the forward-looking (ex ante) risk estimates coming from that model. GEM2L uses a factor structure to model stock returns, which is a popular solution for handling the dimensionality problem inherent to estimating covariance matrices across large numbers of assets. Specifically, a stock's returns are assumed to be driven by a combination of (1) the stock's loadings on systematic risk factors, and (2) by firm-specific "idiosyncratic" effects which are assumed to be uncorrelated across assets. Systematic risk factors include both simple indicator variables (e.g., country and industry membership), along with more complex constructs (e.g., stock-level momentum and value exposures).

The GEM2L risk model estimates factor and idiosyncratic returns via regression using weekly stock-level return data, and then uses an exponential weighting scheme to put more weight on recent data in estimating factor volatilities and correlations. Barra risk models are generally designed to predict risk over a horizon of one

month. However, GEM2L is also built to provide relatively stable risk estimates over time (e.g., it uses longer half-lives in covariance estimation than its shorter-term counterpart GEM2S).

In this paper we focus on three main measures derived from the GEM2L risk model: each stock's total risk, stock-specific (idiosyncratic) risk, and the beta versus the MSCI World index. Total risk captures each stock's overall volatility, whereas stock-specific risk measures the volatility that remains after controlling for systematic risk exposures of that stock. The beta is a measure of the exposure of a stock to the overall market (that exposure being a key component of systematic risk). In Sections 3.1–3.3 we discuss the link between ESG and contemporaneous risk measures. In Section 3.4 we discuss predicting

these risks measures 1, 2, . . . , 5 years ahead using current ESG scores.

3 Do stocks with better ESG exposures have lower risks?

3.1 Summary statistics

We begin with an overview of our data in Exhibit 1. The goal of this table is to introduce the variables we utilize in our subsequent analyses, but also to provide a first glance at the potential relationship between ESG and risk, returns, or firm characteristics and quality profile. To this end, every sample month we sort stocks into quintiles on their ESG scores. For each quintile, we compute the average risk measures of the underlying stocks, average returns, and average firm characteristics. Exhibit 1 presents the time series

Exhibit 1. Summary statistics.

	Q1 (poor ESG)	Q2	Q3	Q4	Q5 (best ESG)	Q5–Q1
Industry-adjusted ESG score	1.5	3.4	4.7	6.2	8.4	6.9 (100.26)
Risk metrics						
Total risk	34.5%	33.1%	33.0%	31.8%	30.4%	−4.1% (−29.35)
Stock-specific risk	24.9%	23.8%	23.7%	22.7%	21.4%	−3.5% (−29.46)
MSCI World beta	1.07	1.04	1.06	1.04	1.04	−0.03 (−11.49)
Performance						
Annualized return	6.52%	4.84%	4.88%	4.85%	4.76%	−1.76% (−1.20)
Characteristics						
Market cap (\$M)	10,364	11,628	13,944	16,895	22,685	12,320 (24.47)
Book-to-price	0.64	0.67	0.65	0.60	0.60	−0.05 (−4.85)
Momentum	10.02%	7.92%	8.70%	8.37%	8.28%	−1.74% (−2.78)
Quality indicators						
Earnings variability	0.55	0.54	0.53	0.51	0.49	−0.06 (−16.96)
Ohlson's credit score	4.29	4.52	4.75	4.55	4.70	0.41 (22.76)
Profitability	0.29	0.30	0.31	0.32	0.33	0.03 (23.49)

Every month, stocks are sorted on their industry-adjusted ESG score into five quintiles. For each quintile, we compute risk measures (total and specific risk, beta versus the MSCI World index), average returns, and firm, and quality characteristics. The table reports the time series averages of these measures, as well as the differences between the two extreme quintiles (stocks with the best versus stocks with the worst ESG profile) and the corresponding *t*-statistics. The sample covers constituent stocks from the Russell 3000, MSCI World ex US, and MSCI Emerging over the period of January 2007 to December 2015.

averages of these quantities, as well as the difference between the two extreme quintiles, that is, the difference between the stocks with the best and the poorest ESG profiles.⁵

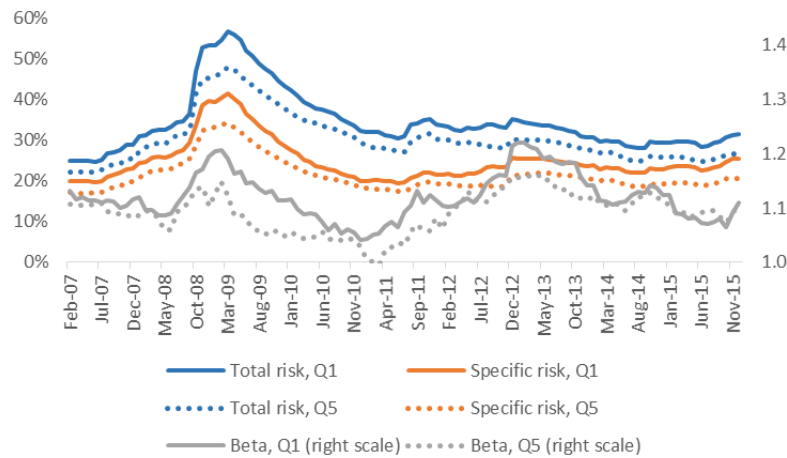
The industry-adjusted ESG score we are using here varies from 0 (indicating large ESG exposures and poor management of those exposures) to 10 (perfect ESG profile). The values are roughly

uniformly distributed across the five quintiles in Exhibit 1.

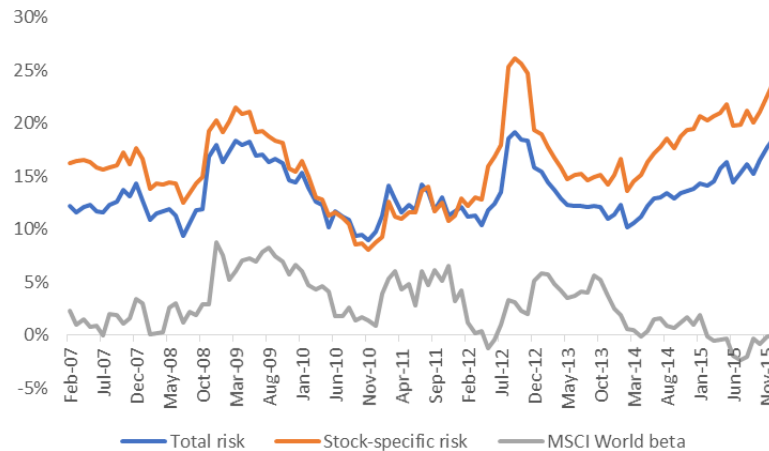
The sort on the ESG score translates into a sort on various risk measures. For example, the total volatility of the average stock in the first quintile (worst ESG) is 35%, versus 30% for the average stock in the fifth quintile (best ESG). We observe similar patterns in stock-specific risk, or in the

Exhibit 2. Top and bottom ESG quintiles: Differences in risk over time.

Panel 1: Total and stock-specific volatility and beta of best (Q5) and worst (Q1) ESG quintile over time.



Panel 2: Percentage change in risk when moving from best to worst ESG quintile.



The first panel presents total volatility, stock-specific volatility, and beta versus MSCI World over time for the worst ESG quintile (Q1) and best ESG quintile (Q5), defined as in Exhibit 1. The second panel shows the percentage change in the three measures of risk when moving from Q5 to Q1, for each sample month. The sample covers constituent stocks from the Russell 3000, MSCI World ex US, and MSCI Emerging over the period of January 2007 to December 2015.

beta of the typical stock: all are higher for poor ESG companies.⁶

Exhibit 2 presents the pattern in the risk measures over time. The average total and stock-specific volatility is higher for the worst ESG quintile (Q1) than for the best ESG quintile (Q5) in every single month of the sample. As the second panel documents, the difference in the risk metrics is quite substantial, with poor ESG stocks having a level of volatility that is higher by about 15% than that of the best ESG stocks. The difference in betas is perhaps less dramatic, but even here for most sample months Q1 stocks have higher betas than Q5 stocks, and the relative difference between the two is about 3%. Overall, the pattern we document seems very consistent over time.

Next, coming back to Exhibit 1, we note that stocks with the worst ESG exposures tend to earn somewhat higher returns. The difference between the two extreme quintiles is of economically large magnitude (approximately 1.8% per annum) but it is not statistically significant, with a *t*-statistic of -1.2 . This is in line with the economic intuition discussed in Section 2: stocks with poor ESG profiles may face less investor demand, leading to relatively lower prices and higher average returns. Put differently, the higher returns might be a premium that investors earn for the displeasure of holdings such stocks, possibly as a compensation for the additional risks these stocks exhibit.

This intuition is supported by the patterns in the stock characteristics we report in Exhibit 1. Stocks with poor ESG scores tend to have smaller size and tend to be cheaper (have higher book-to-price ratios), in line with the idea that the market assigns lower valuations to such companies.

Exhibit 1 also looks at a variety of quality measures, capturing the strength of the company's fundamentals. We incorporate indicators of how

variable earnings are (computed over the trailing five years), Ohlson's *o*-score (a measure of a company's credit risk, with lower values indicating higher creditworthiness, developed in Ohlson, 1980), and gross profitability (gross profit over assets). All these quality indicators may capture risks of a company's fundamentals. We find that stronger ESG profiles are associated with higher quality fundamentals: less variable earnings, better credit risk, and higher profitability.

3.2 Controlling for firm characteristics

We saw from Exhibit 1 that small stocks tend to have poorer ESG scores. At the same time, small stocks tend to have higher risk and higher betas than large stocks do. Thus, the relationship between ESG and risk might be driven indirectly by firm size or perhaps another stock characteristic. To address this point, we turn to regressions that account for firm characteristics and quality profile, as discussed in Section 3.1. We estimate regressions with various risk measures as the dependent variable and the ESG scores as the key explanatory variable. We lag the ESG score by a month: for example, we explain stocks' risks measured by the risk model in February using ESG exposures measured in January of the same year. Formally, the regression model for stock *i* in month *t* is:

$$risk_{i,t} = ESG_{i,t-1} + controls_{i,t-1} + \varepsilon_{i,t}$$

In the simple univariate regressions (1), (4), and (7) we confirm a strong correlation between ESG and statistical risks.⁷ The results are statistically significant, but also relatively large in economic terms. To gauge that, Exhibit 3 reports the estimated impact of a change in ESG from the 75th percentile to the 25th percentile of its distribution. This shift increases total risk and stock-specific risk by over 1.5 percentage points or, relative to the average risk in the top ESG quintile, by about 5%. The corresponding increase in beta is

Exhibit 3. ESG exposures correlate with various dimensions of risk.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Predicting total risk			Predicting stock-specific risk			Predicting beta		
ESG score	-0.433*** (-9.92)	-0.194*** (-4.60)	-0.117*** (-3.23)	-0.465*** (-13.58)	-0.154*** (-5.13)	-0.070** (-2.51)	-0.005** (-2.41)	-0.002 (-0.83)	-0.004*** (-2.85)
Log market cap		-1.739*** (-21.93)	-2.417*** (-34.60)		-2.407*** (-40.25)	-2.779*** (-48.50)		-0.029*** (-7.18)	-0.025*** (-8.98)
Book-to-market		2.127*** (6.76)	1.318*** (5.87)		1.265*** (6.63)	0.905*** (5.58)		0.027*** (3.93)	0.043*** (5.24)
Price momentum		-0.046*** (-16.77)	-0.012*** (-5.71)		-0.032*** (-16.07)	-0.006*** (-3.60)		-0.001*** (-10.44)	-0.001*** (-9.97)
Earnings variability		7.125*** (26.71)	6.338*** (28.53)		4.978*** (25.25)	4.793*** (27.54)		0.294*** (22.59)	0.214*** (24.68)
Profitability		1.171** (2.21)	0.013 (0.02)		3.056*** (7.48)	1.113*** (2.66)		-0.041* (-1.74)	-0.047*** (-2.62)
Ohlson credit score		0.019*** (6.55)	0.025*** (2.65)		0.012*** (6.70)	0.021** (2.32)		0.002 (1.56)	0.001** (2.35)
Observations	263,781	191,188	191,188	263,781	191,188	191,188	263,781	191,187	191,187
R-squared	0.009	0.271	0.618	0.016	0.378	0.604	0.001	0.154	0.571
Fixed effects	NO	NO	YES	NO	NO	YES	NO	NO	YES
If ESG drops from 75th to 25th percentile:									
Change in level of risk	156.2	70.0	42.2	167.8	55.6	25.3	0.02	0.01	0.01
% Increase in risk	5.1%	2.3%	1.4%	7.9%	2.6%	1.2%	1.7%	0.7%	1.4%

Regressions explain various risk measures: total volatility, stock-specific (idiosyncratic) risk, and beta versus the MSCI World market, all obtained from Barra's GEM2L model. Total and stock-specific risk are expressed in percentages. The key explanatory variable is industry-adjusted ESG exposure from MSCI as of the prior month, other variables include company characteristics and quality metrics. Fixed effect, where indicated, add additional controls for the calendar month, GICS sector, and country of domicile for a given stock. *t*-Statistics are based on robust standard errors, double clustered at the firm and the date level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. The table also reports estimated economic impact for a company whose ESG score deteriorates from the 75th percentile to the 25th percentile, both as the change in level and as a percentage change relative to the average risk measures reported in Exhibit 1 for Q5 (stocks with best ESG metrics). The sample covers constituent stocks from the Russell 3000, MSCI World ex US, and MSCI Emerging over the period of January 2007 to December 2015.

somewhat smaller in magnitude, about 2% of the initial level of the beta.

Specifications (2), (5), and (8) incorporate additional controls for firm characteristics and confirm a statistically significant relationship between ESG and statistical risk. The additional explanatory variables help explain considerably more of the cross-sectional variability in risk, as the higher R^2 s indicate. Moreover, some of the explanatory power of ESG scores does seem to be driven

by the correlation between ESG and other stock characteristics, such as firm size. When such characteristics are controlled for, the estimated coefficient on ESG goes down, but it remains economically and statistically meaningful.

Finally, in our most general specifications (3), (6), and (9), we incorporate all these controls, but also add indicator variables for each calendar month in our sample and for each stock's GICS sector and country of domicile. These additional controls are

Exhibit 4. ESG exposures predict risk across various equity universes.

Investment universe	US		World ex US		Emerging	
	Explaining total risk					
	(1)	(3)	(1)	(3)	(1)	(3)
Coefficient on ESG score	-0.348*** (-4.37)	-0.110** (-2.03)	-0.383*** (-5.79)	-0.178*** (-3.18)	-0.477*** (-4.81)	-0.110 (-1.61)
If ESG drops from 75th to 25th percentile:						
Change in level of risk (in bps)	125.5	39.7	138.2	64.2	172.1	39.7
% Increase in risk	4.1%	1.3%	4.6%	2.1%	5.7%	1.3%
	Explaining stock-specific risk					
	(4)	(6)	(4)	(6)	(4)	(6)
Coefficient on ESG score	-0.294*** (-4.31)	-0.046 (-1.06)	-0.402*** (-8.21)	-0.138*** (-3.41)	-0.280*** (-3.86)	-0.138** (-2.57)
If ESG drops from 75th to 25th percentile:						
Change in level of risk (in bps)	106.1	16.6	145.0	49.8	101.0	49.8
% Increase in risk	5.0%	0.8%	6.8%	2.3%	4.7%	2.3%
	Explaining beta					
	(7)	(9)	(7)	(9)	(7)	(9)
Coefficient on ESG score	-0.012*** (-4.56)	-0.006*** (-3.01)	0.012*** (3.55)	-0.003 (-1.37)	-0.001 (-0.28)	-0.007** (-2.36)
If ESG drops from 75th to 25th percentile:						
Change in level of risk	0.04	0.02	-0.04	0.01	0.004	0.03
% Increase in risk	4.2%	2.1%	-4.2%	1.0%	0.3%	2.4%
Additional explanatory variables	NO	YES	NO	YES	NO	YES

We re-estimate key specifications from Exhibit 2 separately for US stocks (Russell 3000 universe), developed stocks outside of the US (MSCI World ex US index universe), and emerging market stocks (MSCI Emerging index universe). Regression numbers (1, 3, etc.) indicate which specification from Exhibit 2 is estimated for a given investment universe. The table also estimates the economic effect of a change in ESG exposure from the 75th to the 25th percentile on the total and specific risk (in bps) and on the MSCI World beta, computed as in Exhibit 3. The sample covers constituent stocks from the Russell 3000, MSCI World ex US, and MSCI Emerging over the period of January 2007 to December 2015.

meant to capture variation in risk over time (some periods are inherently more volatile than others) and variation in risk across sectors and countries (stocks in some sectors, and perhaps some countries, may be more volatile than other stocks). All these additional variables do not change our key finding that ESG is negatively correlated with a given stock's risks.

Regressions in Exhibit 3 establish the overall effect and show that it is robust to controlling for a variety of stock characteristics. To further verify the robustness of our findings we reproduce our main regressions separately for individual markets: US, international developed, and emerging. Exhibit 4 indicates that the pattern is quite universal. Weaker ESG is associated with higher risks in each of these investment universes, with most ESG coefficients having the expected sign and being statistically significant. At the same time, there might be some differences in the magnitude of the effect across markets (e.g., the estimated impact on volatility is lower in the US than in international markets).

3.3 Which components of ESG correlate with risk?

ESG scores aggregate information on exposures to environmental, social and governance issues. So far we found that the combination of such exposures provides information about the risk profile of a given company. A natural question arises here: which dimension of ESG drives this relationship? To investigate this question we look at the environmental, social and governance pillar sub-scores in the IVA data.

In Exhibit 5 we estimate the equivalent of regressions (3), (6), and (9) from Exhibit 3, with total risk, stock-specific risk, and beta regressed on the individual ESG pillars and the usual additional explanatory variables.

We find that across the three dimensions of ESG, it is the social and governance pillars that show the strongest correlation to risk. The environmental pillar is only insignificantly related to the various risk measures. There are two broad reasons why that may be the case. On the one hand, it may be that environmental exposures are inherently less predictive of companies' risks. On the other hand, it is also possible that data on environmental exposures is noisier than data on the other two pillars, and that the noise in the variable prevents the regression from delivering more precise, statistically significant estimates.

3.4 Predicting risks with ESG information

So far, we have documented a strong and robust correlation between ESG exposures and various measures of statistical risk. At a minimum, this shows that ESG exposures reflect some information about risks that is also captured by the risk model. However, is there any additional information captured in ESG scores that is not yet accounted for in traditional risk models? This is a difficult question to answer. The long-run risks we discussed in Section 2 may be difficult to quantify or reliably observe over a fairly short time period for which we have good ESG data. For example, it is probably unrealistic to expect that we could reliably estimate a "rare event" type of risk over the nine years in our sample. This means that we may simply not have enough data to statistically detect potential extreme tail events that ESG data might reflect.

What we can do instead is test whether ESG scores convey information important for future risk, as perceived by the risk model one or more years out. Exhibit 6 presents the results of this analysis.

The dependent variables in these regressions are our usual risk measures: total risk, stock-specific risk, and beta. We relate these measures to ESG scores as reported by MSCI 1, 2, 3, 4, and 5 years

Exhibit 5. Impact of individual components of ESG.

Dependent variable:	(1) Total risk	(2) Stock-specific risk	(3) Beta
Environmental pillar score	−0.046 (−1.03)	−0.039 (−1.14)	0.002 (1.03)
Social pillar score	−0.138*** (−3.24)	−0.102*** (−3.13)	−0.007*** (−3.89)
Governance pillar score	−0.148*** (−5.10)	−0.110*** (−4.60)	−0.004*** (−3.75)
Log market cap	−2.424*** (−29.76)	−2.778*** (−41.81)	−0.026*** (−7.88)
Book-to-market	1.258*** (5.17)	0.862*** (4.89)	0.041*** (3.99)
Price momentum	−0.013** (−2.29)	−0.006* (−1.74)	−0.001*** (−3.78)
Earnings variability	6.291*** (26.37)	4.752*** (25.22)	0.214*** (23.09)
Profitability	0.073 (0.14)	1.170*** (2.74)	−0.046** (−2.41)
Ohlson credit score	0.026*** (2.63)	0.021** (2.31)	0.001** (2.38)
Observations	194,272	194,272	194,271
R-squared	0.619	0.605	0.571
Fixed effects	YES	YES	YES

Risk measures are regressed on E, S, and G components of ESG score and other control variables as in Exhibit 3 specifications (3), (6), and (9).

earlier. Importantly, we also control for the contemporaneous risk measures (e.g., total risk from the Barra risk model, as computed 1, 2, . . . , 5 years earlier) and for the same stock characteristics that we utilized in earlier Exhibits. Formally, the regression model for predicting risk one year ahead:

$$risk_{i,t+1\text{ year}} = ESG_{i,t} + risk_{i,t} + controls_{i,t} + \varepsilon_{i,t+1\text{ year}}$$

ESG exposures help predict future risks as far as three to five years out. As the forecast horizon increases, the magnitude of the estimates goes up as well, consistent with the idea that risks relevant

for ESG exposures may only materialize over the longer term. Statistical significance gets somewhat weaker when we predict risk measures four or five years out, but even here we find the right signs (poorer ESG exposures predict more risk) and significance at least at the 10% significance level.

Importantly, in the regressions we control for the current risk model output (e.g., we predict beta 1 year into the future while controlling for today's beta). Because of that, we interpret the results as indicating that ESG scores convey some information about future risks of the company that may not yet be captured by the risk model (that the

risk model will pick up only a year or two later). Note that the coefficient on the contemporaneous level of risk, while positive and highly statistically significant throughout, declines as our prediction horizon increases. This is to as expected. On the one hand, future risks correlate with current risks (riskier stocks today tend to be riskier in the future as well). On the other hand, the further out we go, the less we can say about a given stock. For example, knowing the beta of a given firm today may be

relatively less important when we predict the beta 5 or 10 years out: eventually, all we can say is that the beta will be just “average,” or say around 1.

While the predictive power of ESG is statistically clear, its economic impact appears modest. As Exhibit 6 reports, a deterioration of ESG from its 75th to its 25th percentile predicts an increase in risk by about 1% of its level over the next few years. At the same time, it would be

Exhibit 6. Better ESG exposures predict lower future risk.

	One year out	Two year out	Three year out	Four year out	Five year out
Predicting total risk:					
Current ESG score	−0.075*** (−3.93)	−0.112*** (−3.28)	−0.102** (−2.46)	−0.114** (−2.27)	−0.130** (−2.26)
Current total risk	0.717*** (47.55)	0.489*** (22.19)	0.426*** (18.55)	0.367*** (14.00)	0.323*** (10.26)
If ESG drops from 75th to 25th percentile:					
Change in level of risk (in bps)	27.1	40.4	36.8	41.1	46.9
% Increase in risk	0.9%	1.3%	1.2%	1.4%	1.5%
Predicting stock-specific risk					
Current ESG score	−0.049*** (−3.13)	−0.069** (−2.56)	−0.072** (−2.16)	−0.079* (−1.91)	−0.097* (−1.95)
Current specific risk	0.693*** (29.68)	0.420*** (13.97)	0.342*** (11.54)	0.322*** (10.59)	0.291*** (8.34)
If ESG drops from 75th to 25th percentile:					
Change in level of risk (in bps)	17.7	24.9	26.0	28.5	35.0
% Increase in risk	0.8%	1.2%	1.2%	1.3%	1.6%
Predicting beta					
Current ESG score	−0.003*** (−3.99)	−0.004*** (−3.49)	−0.004*** (−2.74)	−0.004** (−2.33)	−0.004* (−1.71)
Current beta	0.716*** (48.96)	0.577*** (29.21)	0.548*** (21.17)	0.534*** (17.51)	0.534*** (18.33)
If ESG drops from 75th to 25th percentile:					
Change in level of risk (in bps)	0.01	0.01	0.01	0.01	0.01
% Increase in risk	1.0%	1.4%	1.4%	1.4%	1.4%

Estimation results with risk measures 1, 2, 3, 4, and 5 years out on the current industry-adjusted ESG score, the contemporaneous value of the risk measure, and other control variables as in Exhibit 3 specifications (3), (6), and (9). For brevity, only the coefficients on the ESG score and the contemporaneous level risk are reported here. The table also estimates the economic effect of a change in ESG exposure from the 75th to the 25th percentile on the total and specific risk (in bps) and on the MSCI World beta, computed as in Exhibit 3.

surprising if ESG were a first-order driver of a stock's risk beyond what is already captured in a well-constructed statistical risk model. ESG may convey information about, say, the likelihood of a governance scandal, but such a scandal may not materialize for the average company over a relatively short horizon we study here. Moreover, ESG exposures are fairly persistent and stocks with poor ESG profile tend to be poor ESG also in the future. If we found a very large impact of ESG on future risk, this would imply that the risk of stocks with poor ESG would keep on rapidly increasing over time, which is implausible. So, while it is not clear what a first-principles expectation of the magnitude of the effect should be, the magnitude we find here seems reasonable, leading us to conclude that there is a noticeable relationship between poor ESG profile today and larger statistical risks in the future.

4 Conclusions

In this paper we investigate the relationship between companies' ESG exposures and the statistical risk of their equity. We find a strong positive relationship between the two as stocks with poor ESG exposures tend to have higher total and specific risk and higher betas. We also document suggestive evidence that ESG exposures predict future risks as far as five years into the future. While this pattern is statistically strong, it is of somewhat limited economic importance, possibly because ESG exposures are only noisily measured in the data, or because relatively few ESG risk events materialized during the period we studied. Our relatively short sample period, driven by data availability, also limits our ability to assess the impact of different macroeconomic environments on the strength of these effects.

We interpret these findings as evidence that ESG information may play a role in investment portfolios that goes beyond the ethical considerations and may inform investors about the riskiness of

the securities in a way that is complementary to what is captured by traditional statistical risk models. Investors interested in tilting toward safer stocks may be able to combine the two to build more stable and robust portfolios.

Acknowledgment

We thank Chris Doheny, Andrea Frazzini, Tarun Gupta, Rick Nelson, Chris Palazzolo, Scott Richardson, Kari Sigurdsson and the participants of JOIM Long-run Risks, Returns and ESG Investing Conference for their many insightful comments. The views and opinions expressed are those of the authors and do not necessarily reflect the views of AQR Capital Management, its affiliates, or its employees; do not constitute an offer, solicitation of an offer, or any advice or recommendation, to purchase any securities or other financial instruments, and may not be construed as such.

Notes

- ¹ Much of the existing research has focused on the potential of ESG to influence expected returns. We review such research below and find that it offers a fairly mixed message. For example, if ESG captures a dimension of risk, then one might expect a risk premium from holding stocks with poor ESG profiles. At the same time, some dimensions of ESG (notably, strong governance) have been found to predict positive returns.
- ² For example, the governor of the Bank of England warned about the risks of climate change that a number of firms and institutions may be exposed to ("Carney's warning of carbon's financial risks," the Financial Times, September 30, 2015). Around the same time, University of California quoted risk as a major driver of their planned fossil fuel divestment ("University of California sells coal, oil sands holdings," Pensions and Investments, September 10, 2015).
- ³ For more information please see <http://www.msci.com/products/esg/iva/>. The IVA data are available starting in 1999, but the initial period has poorer coverage and might be structurally different than post-2006 data. For example, in an MSCI whitepaper, Nagy *et al.* (2015) also

use IVA data and write in their footnote 6 that they start their sample in 2007 “to have a homogeneous dataset both in terms of asset coverage and methodological consistency.”

- ⁴ The MSCI ESG data has good coverage for all these universes, at least in the later years of our sample. For example, as of the end of 2015, it covers 97% of the market cap of Russell 1000, 91% of Russell 2000, 96% of MSCI World ex US, and 83% of MSCI Emerging. For some regions the coverage is weaker earlier on in the sample; for example, while the coverage of the MSCI World index has exceeded 90% since 2006, the coverage of MSCI Emerging has exceeded 50% only in 2012.
- ⁵ This technique of averaging cross-sectional estimates is termed Fama–MacBeth after the study that proposed it.
- ⁶ The average beta is above 1 for each of the quintiles in Exhibit 1 because we sample stocks that are not in the MSCI World benchmark (e.g., we include Russell 3000 stocks in the US) and because we compute simple averages within each quintile. The equal-weighted average gives more weight to smaller-cap stocks with higher betas, leading to the numbers in Exhibit 1. When we limit our attention to stocks in MSCI World index, the weighted-average beta of such stocks is of course 1.
- ⁷ The *t*-statistics reported in this and subsequent Exhibits are computed based on robust standard errors, double clustered at both the firm and the date level (i.e., allowing for potential dependence across same-firm observations and same-month observations). The results are robust to other statistical specification, for example Fama–MacBeth.

References

- Bebchuk, L. A., Cohen, A., and Wang Charles, C. Y. (2013). “Learning and the Disappearing Association Between Governance and Returns,” *Journal of Financial Economics* **108**(2), 323–348.
- Borgers, A., Jeroen, D., Kees, K., and Jenke ter, H., (2015). “Do Social Factors Influence Investment Behavior and Performance? Evidence from Mutual Fund Holdings,” *Journal of Banking and Finance* forthcoming.
- Fabozzi, F. J., Ma, K. C., and Oliphant, B. J. (2008). “Sin Stock Returns,” *Journal of Portfolio Management* **35**(1), 82–94.
- Fama, E. and French, K. (2007). “Disagreement, Tastes, and Asset Pricing,” *Journal of Financial Economics* **83**, 667–689.
- Gompers, P., Ishii, J., and Metrick, A. (2003). “Corporate Governance and Equity Prices,” *Quarterly Journal of Economics* **118**(1), 107–156.
- Hong, H. and Kacperczyk, M. (2009). “The Price of Sin: The Effects of Social Norms on Markets,” *Journal of Financial Economics* **93**(1), 15–36.
- Larcker, D. F., Richardson, S. A., and Tuna, I. (2007). “Corporate Governance, Accounting Outcomes, and Organizational Performance,” *The Accounting Review* **82**(4), 963–1008.
- Nagy, Z., Kassam, A., and Lee, L.-E. (2015). “Can ESG Add Alpha? An Analysis of ESG Tilt and Momentum Strategies,” MSCI whitepaper.
- Ohlson, J. A. (1980). “Financial Ratios and the Probabilistic Prediction of Bankruptcy,” *Journal of Accounting Research* **18**(1), 109–131.
- Statman, M. and Glushkov, D. (2009). “The Wages of Social Responsibility,” *Financial Analysts Journal* **65**(4), 33–46.
- US Sustainable, Responsible and Impact Investing Trends, 2014, USSIF, <http://www.ussif.org/content.asp?contentid=82>, accessed on August 1, 2016.

Keywords: ESG; SRI; Environment; Social; Governance; risk