The term structure of carbon premia

Fan Dora Xia and Omar Zulaica[1]

Abstract

This paper explores the carbon premium – the extra yield investors demand to buy bonds issued by firms with more greenhouse gas emissions – in the US corporate bond market. We analyse the carbon premium along two channels, via panel regression. One is the preference channel, under which the premium reflects investors' preference for firms that they perceive as being more environmentally responsible, all else equal. The other is the risk channel, where investors perceive more carbon-intensive firms as being more prone to default. We test the preference channel by investigating the relationship between corporate bond yields and carbon emissions, while controlling for proxies of the probability of default (PD) and for other bond characteristics. We examine the risk channel by analysing how carbon emissions affect the PD. We validate the existence of carbon premia in both channels, with the premium being larger for firms in more energy-intensive sectors. Moreover, the premium differs across maturities, giving rise to a hump-shaped term structure of carbon premia, reaching its highest level at the belly of the curve (maturities of 15–20 years).

JEL classification: G12, G30, Q54.

[1] Bank for International Settlements (BIS). The authors (dora.xia@bis.org, omar.zulaica@bis.org) would like to thank Pierre Cardon, Stijn Claessens, Ingo Fender, Frank Weikai Li, Mark Mason, Michela Scatigna, Hyun Song Shin, Ilhyock Shim, and participants of the BIS Research Meeting,the Banking Department's Reserve Management Seminar and the Seventh Public Investors Conference in Ottawa as well as the tenth ABFER Annual Conference in Singapore for helpful comments; and Adam Cap and Alberto Americo for excellent research assistance. The views expressed are those of the authors and do not necessarily reflect those of the BIS.

1. Introduction

It is an emerging consensus that shifting the global economy to a low-carbon growth path is essential. As more greenhouse gases (GHGs) are accumulated in the atmosphere, global temperatures will continue to rise, with some of the resulting effects being irreversible (IPCC (2022)). Global warming catalyses the frequency and intensity of natural disasters such as droughts and storms. Indeed, in the past three decades, adverse weather events have become more frequent on the back of higher global temperatures. These extreme weather events often cause widespread losses and damage to nature, people and economic activities.[2]

In pricing this transition, investors will likely demand compensation for investing in firms with higher carbon footprints.[3] Why? Investors seem to agree on carbon emissions as a reasonable proxy for gauging the exposure to climate-related transition risk.[4] And they may demand higher yields for financing companies that will be affected by the transition from fossil fuels to renewable energy, for example, thereby giving rise to a *carbon premium*.[5]

Any carbon premium may take into account the following two channels. The first one is the preference channel, reflecting that investors who want to support sustainable growth might have a preference, all else equal, for firms that they perceive as helping to achieve this goal. Seen conversely, investors may dislike firms they perceive as more harmful to the environment. This channel captures aspects of the investment process such as negative screening, which excludes firms that score poorly on environmental factors such as GHG emissions (Elsenhuber and Skenderasi (2019)).

The second is the credit risk channel (the "risk channel"), where investors perceive more carbon-intensive firms as more prone to default.[6] This is because these firms are likely to face larger transition risks, which are related to regulatory policies, advances in technology, and changes in consumer preferences that may impair their financial health. This channel captures practices in banks and credit rating agencies which explicitly take into account environmental factors – such as carbon emissions – in assigning risk grades on their scorecards.

It is important to examine how much of a carbon premium is priced into financial assets for at least two reasons. First, because financial markets can support the transition to a more sustainable economy by reallocating their resources towards economic activities that foster it, it is crucial to evaluate the extent to which this reallocation has affected asset pricing. Naturally, a precondition for this mechanism to work is that investors differentiate between financial assets that fund activities with

2 In addition to transition risk, firms also face physical risk – which directly results from the effects of climate change on economic activities. The threat of disrupted production from rising sea levels to factories located close to the sea is an example. Compared with transition risk, physical risk is harder to quantify. It depends in complex ways on firms' geospatial characteristics.

3 In this paper, we use the term carbon footprint to refer to greenhouse gas (GHG) emissions caused by a firm.

4 For a review, see Giglio et al (2021). That said, one shortcoming of using carbon emissions to quantify exposure to transition risk is the fact that they are backward-looking instead of forward-looking. Some forward-looking indicators are available but do not seem to be widely available/accepted yet.

5 Henceforth we use "carbon premium", "carbon risk premium" and "carbon transition risk premium" interchangeably.

6 We define a *more* carbon-intensive firm as one with *higher* CO_2 emissions than others.

different degrees of environmental impact. On the other hand, climate change can have a substantial impact on financial stability (Bolton et al (2020)). For instance, if carbon risk is not sufficiently priced in, financial assets are vulnerable to sharp re-pricings that could lead to systemic risk episodes. Estimating the carbon premium embedded in current prices sheds light on any risk of such re-pricing.

In this paper, we ask whether such a carbon premium exists in corporate bond prices – specifically, via their yield spread to the risk-free rate. The two channels described above are relatively straightforward to test in this market, given that spreads can be decomposed into a component related to default risk and one capturing all other relevant risk factors, offering us a foundation to dissect the preference and risk angles.[7] To do so, we measure a firm's carbon footprint through its GHG emissions, which we draw from the widely used S&P Global Trucost database. We test the preference channel by analysing the relationship between corporate bond spreads and carbon emissions, while controlling for default risk and other bond characteristics. We test the risk channel by looking into the impact of carbon emissions on the probability of default while filtering out the impact of other firm characteristics. Our analysis focuses on firms in the United States, as it is the jurisdiction best covered by the Trucost database. The market value of US companies covered by Trucost accounts for well over 90% of total US market capitalisation. In those same terms, the firms in the United States account for the largest share (around 40%) of the firms in the database.

Our main finding is that carbon emissions affect the spreads of corporate bonds issued by US firms via both the preference and risk channels.[8] When controlling for the probability of default, we find a positive and statistically significant carbon premium on firm-level scope 1 and scope 2 emissions.[9] We interpret this finding as the *credit risk-adjusted part* of the carbon premium. In a second step, we find a statistically significant and positive relationship between a firm's carbon emissions and its probability of default. We interpret this non-linear relationship as the *credit-risk part* of the carbon premium, which holds for all emission scopes.

Combining both preference and risk channels – credit risk-adjusted *and* credit-risk carbon premia – we can derive *total* carbon premia. For a typical firm in our sample, halving carbon emissions would narrow corporate spreads by around 2 to 4.5 basis points, with the more significant contribution coming from the credit risk-adjusted part of the premium – that is, the preference channel. The larger contribution from investor preferences makes sense at the current juncture. The simplest way to introduce a sustainable investment approach is through screening (ie securities are left outside an investment universe due to their more negative environmental impacts). To the best of our knowledge, frameworks for quantifying the impact of climate events on default risk remain under development.

7 Other advantages of focusing on corporate bonds include that downside risks from climate change are likely to matter more to bond investors than to equity investors. Moreover, investors in corporate bonds are more sophisticated and thus more likely to consider carbon risk. See Duan et al (2021) for details.

8 In addition to our main analysis, we investigate the impact of carbon intensities on spreads through the preference and risk channels, and report these results in Appendix 3. Specifically, we examine how carbon intensities affect corporate bond spreads, and show that these effects are significant through both channels.

9 Scope 1 emissions are direct emissions generated from a firm's activities, while scope 2 emissions are indirect emissions resulting from a firm's purchases of electricity, steam and heating/cooling.

Furthermore, carbon emissions are priced in bonds issued by both non-energy-intensive and energy-intensive firms, with larger impacts for the latter. In particular, through the preference channel, a 50% decrease in the sum of scope 1 and 2 emissions predicts a drop of 8.2 and 4.2 basis points in the spread of bonds issued by energy-intensive firms and non-energy-intensive firms, respectively. For the risk channel, a 50% reduction in either scope 1 or scope 1+2 emissions would reduce the probability of default of a typical firm in an energy-intensive sector by 6 basis points, which can be translated to around 2.4 basis point decline in option-adjusted spreads. The impact for a typical firm in non-energy-intensive sectors is around 2 basis points on the probability of default and around 0.6 basis points on option-adjusted spreads. Putting the two channels together, we find that the combined impact on the bond spread is around 7–13 basis points for firms in energy-intensive sectors, and less than 4 basis points for those in non-energy-intensive sectors. This impact is non-negligible. Taken literally, the result means that, by halving their emissions, firms can improve the credit rating implied by their spread by up to one notch, on average.[10]

More importantly, we find that, for the preference channel, carbon risk loads differently across maturities. The interaction between bond maturity and firm-level emissions is relevant at high levels of statistical significance. We dub this finding the *term structure of credit risk-adjusted carbon premia*. The term structure is hump-shaped. Carbon premia increase with maturity up to the belly of the curve (15- to 20-year maturity) and decline thereafter. We offer two conjectures on the curve's shape. The first is the long-term nature of environmental risks, which, despite requiring critical action today, will become inevitable only in the future. The second is the preferred habitat of investors who operate in this market. For example, institutional investors with a sustainable investment mandate (eg pension funds) may use longer-term bonds to match their liabilities but may not go for the ultra-long segments due to liquidity and interest rate risk considerations. As a consequence, the term structure of total carbon premia is also hump-shaped, because the risk channel introduces (roughly) parallel shifts in the term structure of credit risk-adjusted carbon premia.

To-date, not much is known about whether a carbon premium is reflected in asset prices, and our paper contributes to the small but growing literature on the topic. Bolton and Kacperczyk (2021a, 2021b) focus on the carbon risk premium in equity markets. They document the existence of a widespread carbon risk premium in equities: firms with higher carbon emissions offer higher returns across sectors and countries. Briere and Ramelli (2021) show that a green sentiment index – which captures shifts in investor appetite for environmental responsibility – has explanatory power for stock price performance.[11] Ehlers et al (2022) test whether banks demand a premium when lending to firms with higher carbon emission intensity. They find a statistically significant carbon premium in lending rates across industries in the syndicated loan market since the Paris Agreement was struck in 2015. Huynh and Xia (2021) examine whether climate change news risk are priced in corporate bonds, and find that bonds with a higher climate change news "beta" earn lower future returns. Duan et al (forthcoming) explore the pricing of carbon risk in US corporate bond returns. While we also look at the corporate bond market, we differ from Duan et al (forthcoming) in several aspects. First, we focus on the spread instead of the return because bonds are quoted and traded on this variable. The spread-level angle allows

10 In our sample, the mean spread difference between A– and BBB+ is 10 basis points.

11 The authors, in fact, predict that green sentiment may anticipate a stock outperformance of more environmentally responsible firms.

us to explore carbon risk pricing within and outside default risk. Second, our findings are also different. While Duan et al (forthcoming) conclude that bonds from firms with more carbon emissions offer significantly lower returns, we show evidence of a positive carbon risk premium consistent with what Bolton and Kacperczyk (2021a, 2021b) and Ehlers et al (2022) find in equities and syndicated loans, respectively. At the same time, our observation that both risk and preference channels contribute to a carbon premium is consistent with Briere and Ramelli's (2021) finding that higher investor demand for environmentally responsible stocks is explained by both fundamental and non-fundamental (ie sentiment) motives. Carbone et al (2021) look into how carbon emissions and mitigating measures, such as climate disclosure and emission reduction targets, influence firm credit risk as measured by credit ratings and the distance to default. Consistent with our findings on the risk channel, they also document that high emissions tend to be associated with higher credit risk.

Our paper also contributes to the line of literature investigating whether the environmental and social commitments of firms, more generally, affect their cost of debt. Goss and Roberts (2011) investigate the impact of corporate social responsibility (CSR) performance on the cost of private bank loans and find that banks charge more for loans to firms with social responsibility concerns. Chava (2014) finds a similar relationship between the cost of bank loans and firms' environmental performance. On public debt markets, Ge and Liu (2015) find that firms' better CSR performance is associated with lower spreads after controlling for credit ratings. Polbennikov et al (2016) study the historical relationship between environmental, social and governance (ESG) ratings and corporate bond spreads, finding that bonds with higher ESG ratings have slightly lower spreads, all else equal. More recently, using data from Sustainalytics, Seltzer et al (2022) find that firms' with lower environmental scores tend to have higher yield spreads – carbon emissions being one of the components.

Finally, our paper also adds to the literature on the determinants of corporate spreads. Since at least Collin-Dufresne et al (2001), it has been recognised that spreads on corporate bonds tend to be several times wider than would be implied by expected default losses alone. The phenomenon is widely known as the credit spread puzzle (Amato and Remolona (2003)). To investigate the puzzle, two types of model have been used to estimate corporate spreads: structural and empirical. Our work falls in the latter camp. Empirically, determinants other than default risk such as taxes (Elton et al (2021)), firm-level equity return volatility (Campbell and Taksler (2002)) and liquidity (Chen et al (2007)) have been found useful in explaining US corporate spreads. The above results have also been validated for other markets such as US mortgage securitisations (Fender and Scheicher (2009)). For the euro area, Boss and Scheicher (2002) show that, among other variables, liquidity and equity return volatility are useful in explaining changes in corporate bond spreads. For China, Chen and Jiang (2019) conclude that liquidity risk significantly affects corporate bond pricing, though its contribution is much smaller than its US counterparts.

The debate on the puzzle is very much alive to this day, with papers arguing in favour of or against its existence (see, for example, the contrasting views of Feldhütter and Schaefer (2018) and Bai et al (2020)). We contribute to the ongoing discussion by adding the carbon risk angle. From a specification perspective, our model most closely resembles that of Gilchrist and Zakrajšek (2012), where credit spreads are written as a function of a proxy of default risk and other variables. However, as established above, our primary purpose is not to predict macroeconomic conditions.

The remainder of the paper is structured as follows. Section 2 describes the carbon emissions, and firm-level and bond-level data sets required for our analysis. Section 3 begins the empirical anaysis with the preference channel Section 4 presents the discussion on the risk channel. Section 5 puts together both channels, showcasing the total effect of carbon risk on corporate bonds. Section 6 concludes.

2. Data

Our analysis seeks to explain corporate bond spreads as a function of a firm's carbon footprint, while controlling for bond and firm level characteristics. For this purpose, we collect three types of data: carbon emissions data, firm-level financial data and security-level data, which are matched and merged to produce the database for the analysis. In this section, we describe each of the three types, while pointing to their particular sources.

2.1 Carbon emissions data

Following others in the literature (Bolton and Kacperczyk (2021a, 2021b) and Ehlers et al (2022), for example), we obtain data on carbon emissions from S&P Global Trucost ("Trucost"). The database provides firms' annual carbon emissions for each fiscal year since 2002. Graph 2.1 summarises firm coverage in the data set. The number of firms in the database has expanded from fewer than 2,000 in 2002 to over 20,000 in 2020 (Graph 2.1, red line). Coverage has also broadened in terms of geographic locations, going from firms mainly in advanced economies to firms in both advanced and emerging market economies. In our analysis, we focus on the United States, given that its companies are the lion's share of the database in market value terms. Altogether, these firms account for 40% of total market value of all firms in the database in 2020. These same firms are also approximately 90% of the total US public market capitalisation.

Firm coverage of the S&P Global Trucost emissions database — Graph 2.1
Number of firms

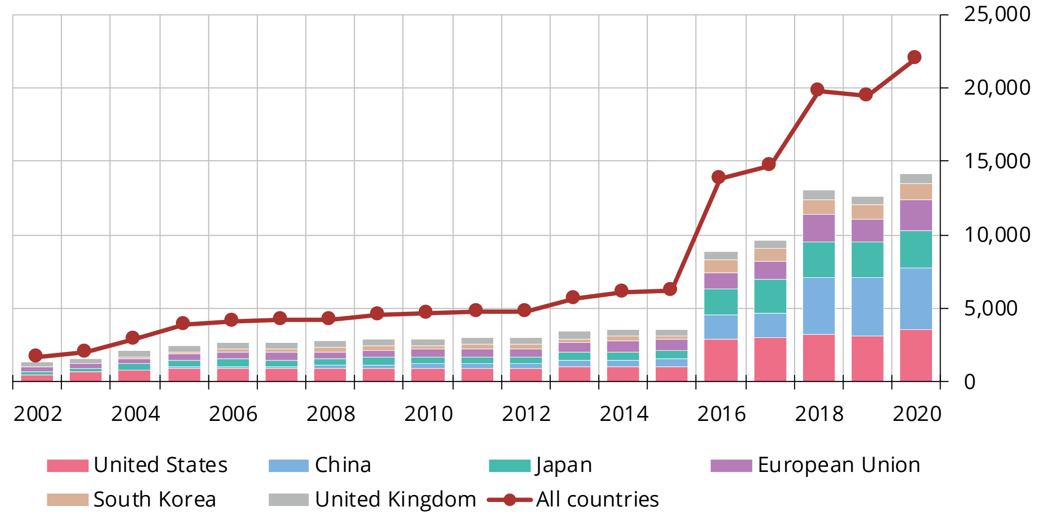

Source: S&P Global Trucost.

Although this database goes back almost two decades, firm composition changed dramatically in fiscal year 2016 due to additions. The red line in Graph 2.1 shows a more-than-twofold jump from 2015 to 2016. Since we do not wish for this change in sample to bias our results, we choose to start our analysis in 2016, where the number of companies is much richer. This also helps keep our carbon emission time series stable over time. Other research has shown that such sample changes can lead to very different conclusions.[12]

We now define the three types of emission that follow the Greenhouse Gas Protocol (GGP) and that are used in our analysis:[13]

- Scope 1 or direct GHG emissions occur from sources that are owned or controlled by the company. For example, emissions from combustion in owned (or controlled) vehicles and emissions from chemical production in owned (or controlled) process equipment. Scope 1 emissions are part of the disclosure requirements in accordance with the GHG Protocol Corporate Standard.
- Scope 2 or indirect GHG emissions. This accounts for emissions coming from the purchased electricity, steam and heating/cooling consumed by the company. According to the GGP, for many firms, purchased electricity represents one of the largest sources of GHG emissions and the most significant opportunity to reduce them. They are also part of disclosure requirements.
- Scope 3 emissions or other indirect GHG emissions. They are a consequence of the activities of a company, but occur from sources not owned by the company (eg extraction of purchased materials and transportation of fuels). The GGP establishes that disclosure of scope 3 is optional, but provides an opportunity to be innovative in GHG management. Given how difficult they can be to measure, the GGP recommends focusing on one or two major GHG-generating activities, instead of performing a life cycle analysis of all products.[14]

To vary our language, we sometimes use the terms "direct", "indirect" and "value-chain" emissions to refer to scopes 1, 2 and 3, respectively. In practice, scope 1 and 2 emissions are widely reported across different data providers, including Trucost. Across providers, these two scopes are highly correlated (by >90%; see Busch et al (2022)), which is a sign of consistency. However, this is not the case for scope 3 emissions, given the optional nature of their reporting. As a consequence, one needs to estimate them and methodologies vary across suppliers (eg Trucost uses an input-output method). Given that the data quality of scope 3 emissions is questionable,

12 In their work, Bolton and Kacperczyk (2021a) explain that the important shift on average carbon emissions from 2015 to 2016 is due to the inclusion of new firms. When analysing the effects before and after the Paris Agreement, they also find that excluding these new firms, the carbon premium they find with the full sample becomes statistically insignificant.

13 The Greenhouse Gas Protocol establishes comprehensive global standardised frameworks to measure and manage greenhouse gas emissions from private and public sector operations, value chains and mitigation actions.

14 See GGP (2020), Chapter 4 "Setting Operational Boundaries". Scope 3 emissions include both upstream and downstream emissions. In our analysis, we focus on upstream emissions, as they are relatively easier to estimate and therefore have longer time series available.

anecdotal evidence suggests only a very limited use of this measure in investment decision-making.

Carbon emissions vary over time and across firms

Graph 2.2

Average carbon emissions in the US by fiscal year

Thousands of tonnes of CO_2

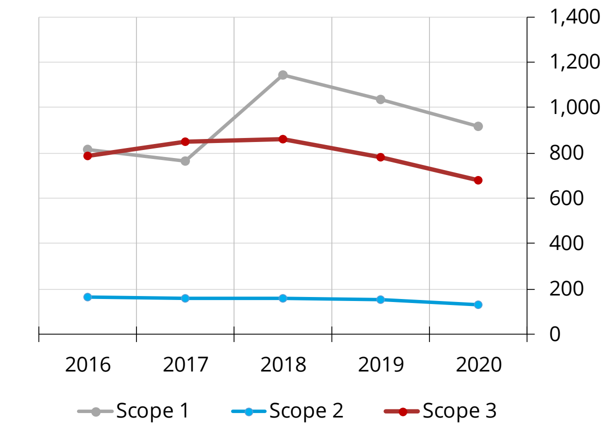

Distribution of US log-carbon emissions in 2020[1]

Probability density

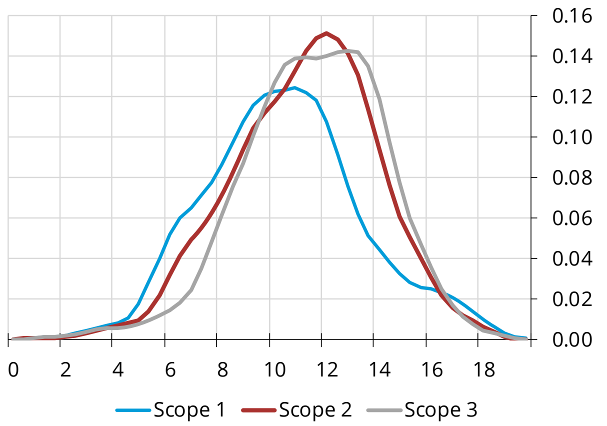

[1] The distribution of carbon emissions is highly skewed. Therefore, a natural logarithm transformation is applied before building the kernels.

Sources: S&P Global Trucost; authors' calculations.

The left-hand panel of Graph 2.2 plots average carbon emissions by scope for all US firms in the Trucost database. Analysing their magnitudes, we see that scope 1 emissions appear to be the highest at about 920 thousand tonnes of CO_2 in 2020; this is followed by scope 3 (~680 ktCO_2e), and scope 2 (~130 ktCO_2e). We also observe that average emissions declined by some 20% from 2018 to 2020, probably reflecting corporations' efforts to reduce their share of GHG. However, focusing on the mean masks an important fact: emissions have wide and asymmetrical distributions. The right-hand panel of Graph 2.2, illustrates this point. Once we take the natural logarithm of firm-level emissions (which are skewed), we see wide and rich probability density functions. Surprisingly, all appear rather continuous and any skew left is not overly pronounced.[15] The panel also depicts the higher mean of scope 3 (yellow line) emissions, although they are defined in a tighter range, than that of scope 1 emissions (red line), for instance.

Across sectors too, carbon emissions vary substantially. Graph 2.3 shows average emissions in the United States by GICS sector in 2020.[16] For direct emissions (red bars), sectors traditionally perceived as brown such as utilities, energy and materials stand out as the top three. Yet, when indirect (blue bars) and value-chain (green bars) emissions are taken into account, consumer staples becomes one of the top carbon-intensive sectors. Based on this, and for the purpose of our analysis, we classify utilities, energy and materials as "energy-intensive" sectors when we talk about scope

[15] The probability densities of carbon emissions in tonnes of CO_2 are highly skewed, and therefore require an adjustment before being used in a regression model. The brief exercise illustrates the case of applying a log transformation.

[16] GICS stands for the Global Industry Classification Standard – a method for assigning companies to a specific economic sector and industry group that best defines its business operations. It consists of 11 sectors.

1 emissions and the sum of scopes 1 and 2. And, we replace materials with consumer staples as part of the energy-intensive category when we include scope 3 in our firm-level total emissions.[17] It is also interesting to analyse how total emissions are distributed *within* sectors. When computing the share each scope represents as a percentage of total, we find that scope 3 makes up for a great share in many sectors. Focusing only on emissions with disclosure requirements, however, the results are mixed. Depending on the sector, total emissions at the firm level may be driven by their direct (eg in the energy sector) or indirect emissions (eg in real estate), which is certainly dependent on the nature of the business. We keep this in mind when we consider the existence of the carbon premium.

Average carbon emissions by sector in 2020 — Graph 2.3

In thousands of tonnes CO_2e

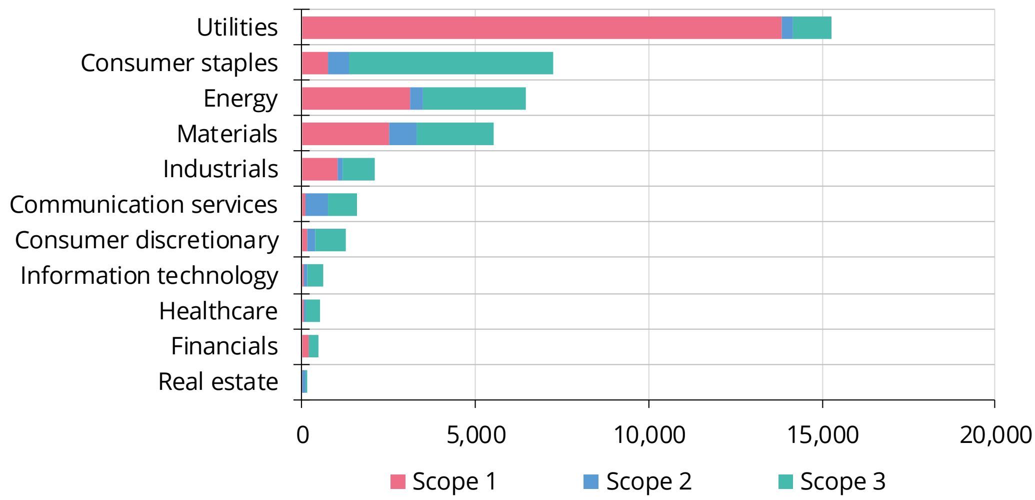

Sources: S&P Global Trucost; authors' calculations.

2.2 Firm-level data

In addition to carbon emissions data, we make use of two other types of firm-level data. The first is a measure of credit risk: the probability of default, our key variable to examine the risk channel. The second includes other firm characteristics affecting their credit risk that are well established in the literature. We need to control for these variables when testing whether a firm's carbon emissions play a role in its credit risk. We gather these variables for companies with carbon emissions data.

2.2.1 Probability of default

To measure a firm's credit risk, we rely largely on default probability data from Bloomberg. Bloomberg provides forward-looking real-world probabilities of default for publicly traded firms. As these are updated daily, the estimates are up to date with current market conditions. A logistic model is used to estimate the probability of

[17] A dummy variable distinguishing between "energy-intensive" and "non-energy-intensive" firms is used in our analysis.

default based on factors that best capture credit risk.[18] For firms in our sample, the annualised five-year-ahead probability of default ranges from near zero to around 32%, with mass concentration being in the range of 0–4% (Graph 2.4, left-hand panel).

For our robustness checks (in section 6), we also compute our own probabilities of default, which we derive from the Merton (1974) model. Details of the computations can be found in Appendix 1. The two measures of default probability are plotted in Graph 2.4 (right-hand panel). The default probabilities are different because our estimates consist of risk-neutral default probabilities, while Bloomberg's estimates are physical or "real-world" probabilities. The difference between the two reflects risk premia.[19] While the latter is more relevant for corporate bond pricing, the former has the advantage of being a cleaner measure of credit risk.

Five-year-ahead probability of default

Graph 2.4

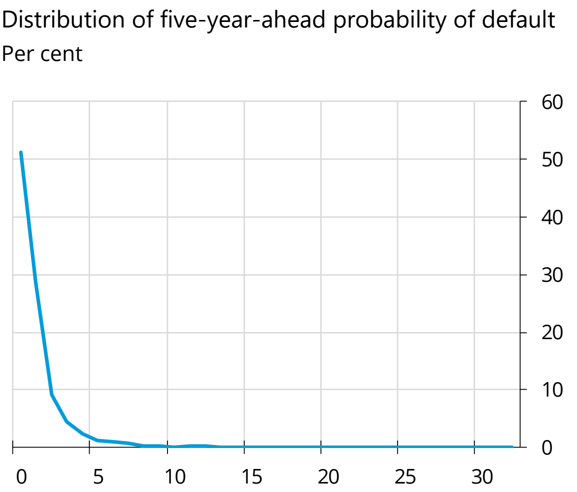

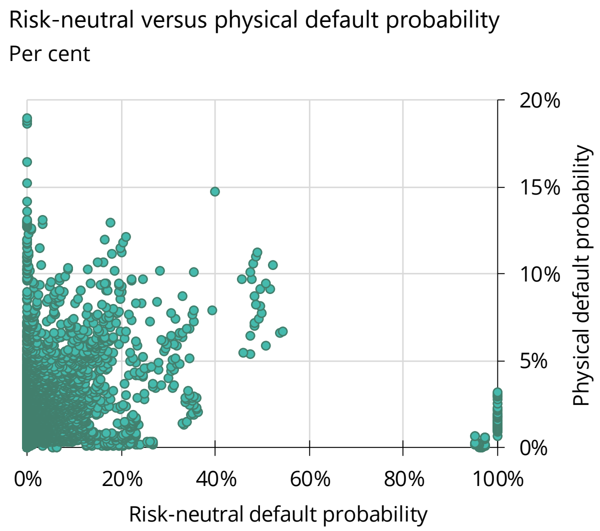

Sources: Bloomberg; S&P Capital IQ; authors' calculations.

An alternative measure of credit risk is the credit rating provided by agencies such as S&P Global, Moody's and Fitch.[20] We prefer the probability of default to credit ratings because the latter are coarsely grained, with the rating designed to remain broadly static over time. Probabilities of default fluctuate over the short term, reflecting information at a higher frequency. We further analyse firm-level default probability when we look at the security-level data in section 2.3.

18 The risk factors include relevant accounting ratios such as return-on-assets, non-performing loans for financial firms, and interest coverage for non-financial firms as well as the distance to default. The distance to default is calculated using the Black-Cox model, which writes it as a function of total debt (proxied by the sum of short-term debt and 50% of long-term debt), value of assets and the implied volatility of assets.

19 Under the Merton model, default takes place in a contingent claims framework, which means all probabilities are risk-neutral. Bloomberg's model (see Bondioli et al (2021)) adds an extra step, where risk-neutral probabilities are mapped into physical ("real-world") probabilities via a logistic model. Nonetheless, under the structural modelling framework, it is easy to show that the spread can be written as a function of the risk-neutral default probability, giving us permission to use this quantity as a regressor.

20 For example, Carbone et al (2021) examine how a firm's carbon footprint affects its credit rating.

2.2.2 Financial variables on firms

To control for variables that affect a company's credit risk, we obtain relevant firm-level financial variables from S&P Capital IQ.[21] The control variables we consider include: size of assets, the long-term debt-to-asset ratio, the earnings-to-asset ratio, the capital-to-asset ratio and return-on-assets. We choose these variables to be consistent with the literature on corporate credit risk (eg Carbone et al (2021)). Intuitively, larger, more profitable and better-capitalised firms are typically less likely to default. By contrast, more indebted firms are more prone to default. Summary statistics for these variables are provided in Table 2.1. Note that these variables tend to be defined in wide ranges relative to their respective standard deviations. Therefore, we winsorise them at the 2.5 percentile before conducting our regression analysis.

We also obtain daily equity price data from Bloomberg. We use these data to compute the volatility of equity returns, a firm-level characteristic that we will control for in the preference channel analysis.

Summary statistics for financial variables on firms

Table 2.1

	Mean	Std. dev.	Min	Median	Max
ln(asset)	7.61	1.96	–4.26	7.62	15.14
Long term debt/asset	0.29	0.52	0	0.24	30.26
Earnings/asset	–0.86	22.72	–1,312.97	0.05	7.65
Capital/asset	0.09	1.8	–119.33	0.13	1
Return on asset (%)	0.39	22.43	–3,487.93	2.45	677.93

Summary statistics are computed across 2,813 firms between January 2017 and December 2021.

Sources: S&P Capital IQ; authors' calculations.

2.3 Security-level data

To build our bond data panel, we start from the universe of all firms with emissions data from 2016 to 2020 in the S&P Global Trucost database. We conduct our data-gathering process in two steps: first, we make a list with the corporate bonds of firms with carbon emissions data; then, we fetch the relevant data fields for this list of securities.

To find the securities issued by each firm in our carbon data set, we use Refinitiv. Its search function allows us to use company-level ISINs (found in the Trucost database) to generate individual CUSIP lists for each of the firms. We include bonds issued by both the parent company and its subsidiaries. To query these lists, we apply a series of filters, in the spirit of Bai et al (2019). We exclude the following:

- Bonds that are not listed or traded in the US public market, which includes bonds issues through private placement and bonds issued under the 144A rule;[22]
- Structured notes, asset-backed, equity- or index-linked securities;

[21] Macroeconomic variables affect a firm's credit risk as well. We include time-fixed effects to this end.

[22] Unlike Bai et al (2019), we preserve bonds traded in a currency other than the US dollar.

- Convertible bonds;
- Floating coupon rate securities;
- Securities with a maturity lower than one year;
- Unrated securities; and
- Bonds trading under $5 or above $1,000.

For the remaining securities, we download a set of static and historical data fields by combining two sources: Bloomberg and Refinitiv. And, as we will be using emissions data lagged by a year (which start in 2016) in our model, these historical fields are obtained from January 2017 to December 2021.

From Bloomberg, where trade data are more widely available, we fetch monthly option-adjusted spreads and daily close prices. The latter will be used to compute our time-varying measure of bond liquidity (see Annex 1 for details) – an important determinant of corporate bond spreads.[23] From Refinitiv, we download monthly data for the amount outstanding, maturity, age, duration and credit rating of each bond. From this same source, we draw static data on coupon and whether the security is callable or not, which we store as a dummy variable (equal to one if the bond is callable and zero otherwise).

We then perform some transformations. First, to ensure that our results are not driven by a small number of extreme observations, we winsorise option-adjusted spread data at the 2.5% fraction. Next, we assign a numerical score to credit ratings. Our score goes from 0 to 20, where the AAA rating on the S&P scale is assigned the highest value (20) and a rating of C is the lowest (zero). This way, we ensure that our variable has the following interpretation: higher credit quality entails a higher credit score.

Summary statistics for corporate bonds — Table 2.2

Variable	Mean	Std. dev.	Min	Median	Max
Option-adjusted spread[1] (bps)	141	114	17	111	619
Maturity (years)	10.4	10	1	6.8	101.4
Duration (years)	7.1	5.1	0	5.7	35
Age (years)	5.3	5.6	0	3.7	32.6
Coupon (pp)	4.2	1.8	0	4	12.3
Amount outstanding (USD mn)	1,050	5,000	0	535	250,000
Credit rating[2]	13	3	0	13	20
Liquidity[3] (bps)	27	27	0	18	122
Callable (binary)	62%	-	0	-	1

Sample period: January 2017 to December 2021; observations = 263,797; number of bonds: 7,599. 1 Option-adjusted spreads are winsorised at the 2.5% fraction. 2 Credit ratings are converted to a numerical S&P scale equivalent from 1 to 20, where 20 = AAA+ and 1 = CCC. 3 Absolute roll measure in basis points (see Annex 1).

Sources: Bloomberg; Refinitiv; authors' calculations.

Our final sample comprises 7,599 securities issued by 779 unique firms. Table 2.2 shows the set of cross-sectional summary statistics we use to characterise our bond

[23] In fact, our full list of determinants is based on the corporate bond literature. This is covered in depth in section 3.

sample. We draw two important observations from the table. First, option adjusted spreads are straddled in a wide range from 17 to 619 basis points, with the average being about 140 basis points. Second, the average credit rating is 13 (equivalent to BBB+), two notches above the investment grade cutoff. With regards to other features, we see that, by construction, bond maturity is above 12 months and is 10 years on average. In line with its definition, modified duration stands somewhat lower, at 7.1 years. We also see the average coupon sitting at about 4 percentage points per annum, and the outstanding issue size slightly above $1 billion on average. It is also important to note that over 60% of our sample is constituted by callable bonds, which asserts our choice of *option-adjusted* spreads to account for this optionality.

We can also look at our bond data under different sample splits. Table 2.3 shows the average corporate spread across the different GICS sectors, which we know are important when looking at firm-level GHG emissions. Across industries, the mean spread sits between 100 and 250 basis points, with the lowest in information technology (103 basis points) and the highest in the energy sector (242 basis points). In theory, this heterogeneity reflects the differences in credit risk across sectors, which underscores the need to control for firm-level default probability. Indeed, this ordinal relationship is preserved when we look at firm-level default probability (fourth column in Table 2.3). Furthermore, spread dispersion, captured by the standard deviation, appears to differ from one sector to another, highlighting the nuances *within* sectors. When looking at bond maturity, we find that the sector average is close to the full-sample number of 10 years in most cases. An exception appears to be the real estate sector, with an average maturity of seven years. Finally, when we look at the number of observations to be included in the model, we see a slight dominance of the financial and industrial sectors. The industry categories with fewer observations are materials, utilities and real estate.

Corporate bond summary statistics by sector

Table 2.3

Sector	Mean spread[1]	Std. dev. of spread	Firm's probability of default[2]	Bond maturity	Number of observations
Communication services	155	102	0.50%	13	21,723
Consumer discretionary	182	140	0.70%	9	25,819
Consumer staples	108	91	0.30%	10	27,914
Energy	242	176	1.20%	10	17,782
Financials	126	92	0.40%	9	40,693
Healthcare	108	81	0.30%	11	31,322
Industrials	138	108	0.60%	12	34,735
Information technology	103	82	0.30%	10	28,551
Materials	174	120	0.50%	10	12,402
Real estate	166	116	0.40%	7	10,910
Utilities	126	73	0.50%	10	11,831
All sectors	141	114	0.50%	10	263,682

Sample period: January 2017 to December 2021; observations = 263,797; number of bonds: 7,599. [1] Option-adjusted spreads are winsorised at the 2.5% fraction. [2] Annualised, five-year-ahead probability of default from Bloomberg.

Sources: Bloomberg; Refinitiv; authors' calculations.

Finally, exploring the data by rating (table omitted for brevity), we find a strictly monotonic relationship between credit rating and corporate spreads: a higher credit quality translates into a lower mean spread. For instance, the mean spread on AAA–rated securities is 54 basis points, with the spread tripling at the investment grade cutoff of BBB–. The result validates the magnitude of the spread as the market's proxy for the perceived level of default risk. Naturally, credit ratings are only a qualtitative (or "soft") indicator, and a model is required to quantify default risk. Indeed, we again present the average value of the firm-level probability of default (fourth column of Table 2.4). The relationship of this variable with the mean spread is almost strictly monotonic, except for the break in the BB– notch, where the average default probability is 0.94%. Nonetheless, these differences are explained by the reduction in sample size as we approach the lowest credit ratings. Indeed, these are the least represented in our sample: altogether, credit ratings of BB+ or below represent about 14% of all our observations, which means our bond sample best (yet not exclusively) represents the investment grade corporate debt spectrum.

Putting the above findings altogether, we conclude that it is important to control for sector-, firm- and issue-level features when performing our regression analysis, which is covered in Sections 3 and 4. Section 3 is dedicated to the analysis at the corporate spread level – the preference channel – and Section 4 details the default risk probability model – the risk channel – which complements our headline results.

3. The preference channel

In this section, we establish a model which, controlling for default risk, is able to explain over 80% of the variation of credit spreads in the United States. We then extend it to include carbon transition risk. Our hypothesis is that, default risk considerations aside, investors trade on information about firms' carbon emissions – a gauge of their environmental footprint and, thereby, of their exposure to carbon transition risk. The consideration of carbon transition risk beyond credit risk captures a genuine preference on the part of investors for environmentally friendly firms, whether motivated by the firms' reputations or investor mandates. Such preferences may be reflected in practice, for example, in the screening of issuers when building investment portfolios. Whether default risk *itself* is affected by the carbon footprint of a firm is addressed in Section 4.

3.1 Underlying theory

According to theory, the price of a corporate bond must reflect the spot rate of a default-free bond (ie government bond) plus a risk premium paid for facing default risk and any options embedded in the issue. This risk premium is known as the corporate spread, and is computed as the difference between the risk-free rate and the yield to maturity on the corporate bond. We denote the spread of bond j issued by firm i at time t as $s_{i,j,t}$, and the firm's probability of default with $P_{i,t}$.

Our empirical methodology is based on the premise that the spread on a bond is directly proportional to the issuing firm's probability of default ($\beta > 0$ a constant):

$$s_{i,j,t} = \beta P_{i,t} \qquad (1)$$

The higher $s_{i,j,t}$, the greater the expectation that the firm will fail on its payments. We can validate this by taking a look at the relationship between these variables in practice. The left-hand panel of Graph 3.1 shows, for January 2021, firm-level default probabilities in the United States (proxied here by the five-year probability described in Section 2) plotted against the option-adjusted spreads on a firm's bonds.[24]

The relationship appears positive: the greater a firm's default probability, the wider the spread on its securities. There are some nuances, however. For example, looking at firm-specific information ignores important bond features, such as the maturity of the issue and its particular credit rating. It is well known that credit spreads increase with respect to credit ratings, which are qualitative assessments about a bond's serviceability. The right-hand panel of Graph 3.1 illustrates how the average spread increases as a function of these two variables: maturity (going from left to right) and rating (going from bottom to top). The function linking maturity with spread is known as the *term structure* of credit spreads.

US credit spreads and their relationships to firm- and bond- specific variables — Graph 3.1

Relationship between a firm's default probability and its bond credit spread (as of Jan 2021)[1]
Percentage; basis points

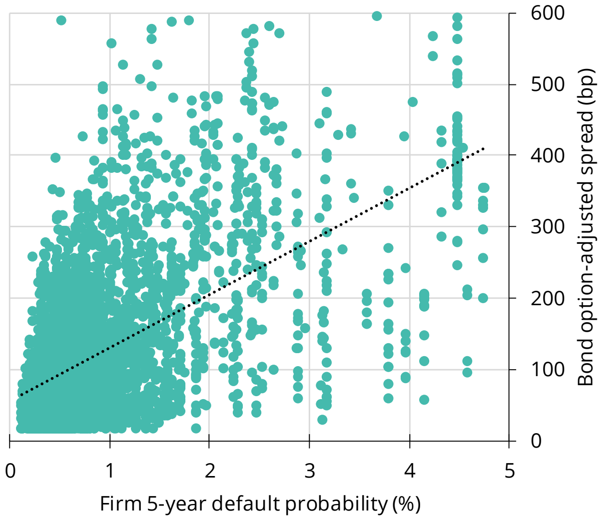

Term structure of credit spreads, by bond rating and maturity buckets (as of Jan 2021)[2]
Basis points

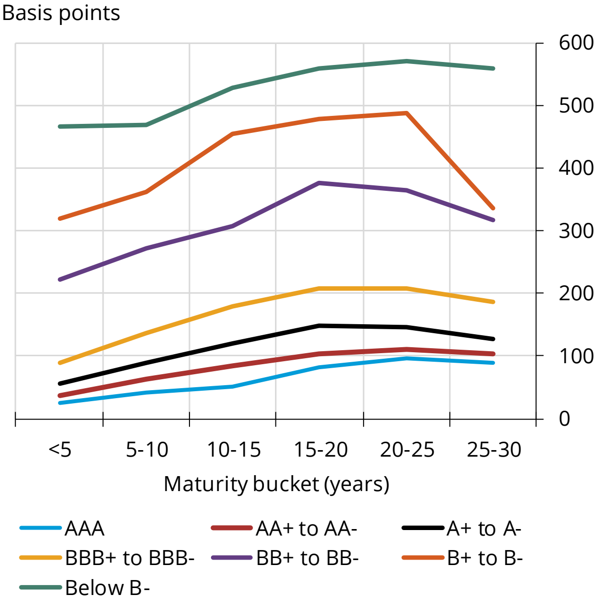

[1] Dashed line denotes a simple regression line $ax + b$. [2] Computed as the average credit spread across all bonds within each combination of maturity and credit rating buckets.

Sources: Bloomberg; Refinitiv; authors' calculations.

The fact that the default probability alone cannot explain all of the variation in spreads is an established feature from the literature, known as the credit spread puzzle.[25] More specifically, the puzzle derives from the fact that neither levels nor changes in the yield spread of corporate bonds over Treasury bonds can be fully explained by credit risk determinants proposed by structural models (eg a firm's

24 Given that the credit risk premium can reflect any options embedded in the issue, we use the option-adjusted spread in our analysis. This is important because, as explained in Section 2, above 60% of our sample comprises callable bonds.

25 This is an empirical finding held since at least Collin-Dufresne et al (2001). For a much more recent discussion on the credit spread puzzle, see Bai et al (2020).

financial health, macroeconomic conditions). As a consequence, the model for spreads is much more complex than in equation (1), looking more like equation (2) below:[26]

$$s_{i,j,t} = \alpha + \beta_P P_{i,t} + \beta_Z Z_{i,j,t} + FE + \epsilon_{i,j,t} \quad (2)$$

where α is the constant of the regression model; β_P the coefficient on the firm-specific probability of default $P_{i,t}$; $Z_{i,j,t}$ a vector of bond- and/or firm-specific controls (with β_Z their respective coefficients); FE a set of fixed effects (typically related to the macroeconomic environment and firm-specific characteristics) and $\epsilon_{i,j,t}$ a zero-mean disturbance or "pricing error". A body of research has been dedicated to defining and expanding the set of controls $Z_{i,t}$, adding missing pieces to the puzzle. We review these next, in chronological order:

- Elton et al (2001) propose the tax premium as a determinant of yield spreads. According to their work, this premium arises because of the higher taxation of corporate bonds compared with sovereign bonds. This effect was later debated by Longstaff (2005), who finds weak support for the hypothesis that the non-default component of spreads is due to taxes.
- Campbell and Taksler (2002) find that idiosyncratic firm-level equity volatility is directly related to the cost of borrowing for corporate issuers. According to them, volatility should drive up the yields of bonds, given that volatility of firm value hurts bondholders. Their study suggests that volatility can indeed explain cross-sectional variation in yields as much as credit ratings. This result has been carried forward to more recent models, such as Rossi (2014), who works with realised volatility.
- Chen et al (2007) argue in favour of the existence of a liquidity premium. They show that several measures of corporate bond liquidity such as the bid-ask spread or the percentage of zero returns are key determinants of bond yield spreads. A wide array of studies has included liquidity as a standard variable in corporate bond modelling; a more recent example is He and Milbradt (2014).

Our research adds to this list by considering a measure capturing a firm's carbon emissions. Our hypothesis is that carbon transition risk is priced in the cross section of corporate spreads, thereby granting investors a carbon premium. As in the case of the liquidity premium, the risk arises from holding a bond that is *less preferred* by investors, given a firm's heavier environmental footprint relative to others. As argued in Ehlers et al (2020), environmental factors – and, most importantly, carbon emissions – are a material financial risk for creditors, which invites exploration.[27]

3.2 Baseline estimation

Our exercise consists of estimating equation (2) via panel regression at the bond level. We estimate $s_{i,j,t}$, the spread of the bond j issued by firm i at month t, as a function of $P_{i,t}$, the firm-level estimate of five-year-ahead default probability at month t, plus the following other variables:

[26] See Gilchrist and Zakrajsek (2012) for example, where option-adjusted spreads are a function of distance to default (a variable representing default), plus bond- and firm-specific variables.

[27] This section covers the preference angle, however. The credit risk angle is handled separately in Section 4.

- $Z_{i,j,t}$ is a vector of six bond-specific variables and one firm-level variable. On the bond side, we use duration, age, coupon and (the natural logarithm of) the amount outstanding as controls. Because liquidity is a well known determinant, we also compute the Roll measure of liquidity, which serves as our proxy for bid-ask spreads.[28] Furthermore, we include a dummy variable which is equal to 1 when the issue is callable and zero otherwise.[29] For the firm side, we compute company-level equity return volatility, as in Campbell and Taksler (2002).[30]
- FE, is a battery of fixed effects. Time-, firm- and credit rating- fixed effects are included.[31] As in the literature, time fixed effects serve as controls for macroeconomic effects (eg state of the yield curve, business cycle). The latter, as in Gilchrist and Zakrajšek (2012), are meant to capture the "soft information" regarding the firm's financial health, which is complementary to our default probability measure.
- Finally, to test our key hypothesis: that carbon transition risk is priced in corporate bond spreads, we need a metric of carbon emissions.

Which measure(s) should be included in the model? We take a practitioner's view and assume that, when making their decisions, investors care about whether the company pollutes the environment or not, regardless of profit.[32] We also suppose that, when they look at greenhouse gases, they think of them on a cumulative basis.[33] In other words: investors do not consider indirect emissions (scope 2) independently of direct emissions (scope 1, the baseline). Instead, they care about the total level of pollution: the sum of both scopes altogether. It is also important to note that the reliability of value chain emissions (scope 3) is at this stage questionable, given their lack of wide availability and inconsistency across data providers.[34] We keep this in mind when analysing our results.

28 This measure of illiquidity was originally proposed by Roll (1984). More recently, Christopoulos (2021) introduced a version which addresses the presence of positive autovariance in the original formula. See Appendix 2 for computational details.

29 Duffee (1998) finds that the relation between credit spreads and Treasury rates is stronger for callable bonds than for non-callable bonds. It is thereby important to make a distinction between these two types of instruments in any spread model.

30 As in Campbell and Taksler (2002), equity return volatility is calculated as the 180-day trailing standard deviation of the firm's stock return at the end of each month.

31 The composite credit rating is the average rating across three providers: S&P, Moody's and Fitch, when available. Furthermore, we assume that, if present, the effects of taxes are absorbed by the fixed effects, given their static nature.

32 It is debated whether the explanatory variable representing emissions should be a level or a ratio (eg intensities). We pose that investors care whether a firm is pollutes *more* or *less* than others, regardless of their level of profitability. This is because, when the ultimate goal is "net-zero", firms who emit more greenhouse gases into the atmosphere are not less exposed to a carbon tax, technological change or investor dispreference simply because they generate more income.

33 Nonetheless, we investigate the impact of carbon intensities (defined as the ratio of carbon emissions to revenue) on spreads and report these results in Appendix 3.

34 Scope 1 and scope 2 emissions have been more systematically reported because of disclosure requirements. However, scope 3 emissions are still estimated by data providers, such as Trucost. Busch et al (2022) find that the complexity of carbon accounting increases from direct emissions to indirect emissions (scope 2 and 3), and the consistency of data between third-party providers decrease: correlations among them drop from >90% to <60% across providers. They suggest that

Therefore, in our regression, we include a firm-level term capturing total GHG emissions in tonnes for each financial year. We work, first, with scope 1 emissions; then, with the sum of scopes 1 and 2; and finally, with the sum of scopes 1, 2 and 3. As others in the literature, we lag these numbers by 12 months, to reflect the availability of this information for the average investor.[35] For easier interpretation of the coefficient, we take their natural logarithm.

Our estimated model is as follows:

$$E(s_{i,j,t}) = \hat{\alpha} + \hat{\beta}_P P_{i,t} + \hat{\beta}_Z Z_{i,j,t} + \hat{\beta}_{P,Carbon} \ln(\text{Emissions}_{i,t-12}) + \widehat{FE} \qquad (3)$$

The terms following $\hat{\beta}_P P_{i,t}$ in the equation represent spread determinants beyond credit risk. This allows us to conjecture that the effect captured by our estimate $\hat{\beta}_{P,Carbon}$ is due to investor preferences, all else equal. We call the effect of this coefficient the preference channel.

We present the results from four regressions in Table 3.1. The first is a specification without carbon emissions and the latter three introduce emission scopes 1 to 3 in a cumulative fashion. We start by focusing on specification (1) to analyse the effect of well known bond- and firm-level controls.

Effects of carbon emissions on US corporate bond spreads

Table 3.1

	(1)	(2)	(3)	(4)
ln(scope 1 emissions)		1.61**		
		[0.74]		
ln(scope 1+2 emissions)			5.22***	
			[1.16]	
ln(scope 1+2+3 emissions)				2.44
				[1.75]
Default probability (%)	32.01***	31.99***	31.95***	32.04***
	[1.26]	[1.26]	[1.26]	[1.26]
Duration	5.22***	5.22***	5.22***	5.22***
	[0.12]	[0.12]	[0.12]	[0.12]
Age	0.56***	0.56***	0.56***	0.56***
	[0.13]	[0.14]	[0.14]	[0.14]
Coupon	10.53***	10.54***	10.53***	10.52***
	[0.47]	[0.47]	[0.47]	[0.47]
ln(amount outstanding)	-2.92***	-2.92***	-2.92***	-2.92***
	[0.32]	[0.32]	[0.32]	[0.32]
Equity return volatility (%)	17.78***	17.74***	17.73***	17.79***
	[0.93]	[0.93]	[0.93]	[0.93]
Liquidity	0.44***	0.43***	0.43***	0.43***
	[0.17]	[0.17]	[0.17]	[0.17]
Callable	-8.00***	-7.97***	-7.95***	-7.98***
	[0.99]	[0.99]	[0.99]	[0.99]
Number of bonds	7599	7599	7599	7599
Observations	263,682	263,682	263,768	263,797
R-squared	0.84	0.84	0.84	0.84

*** p<0.01, ** p<0.05, * p<0.1. Specifications with time-, firm- and credit rating fixed effects. Standard errors clustered at the security level.

Sources: Bloomberg; Refinitiv; authors' calculations.

requesting firms to follow a standardised approach is even more important in the complex scope 3 realm.

35 See eg Ehlers et al (2022) and Duan et al (forthcoming).

The probability of default is, as expected, highly significant in explaining corporate bond spreads. Their magnitude is also powerful: an increase of 1 percentage point in this probability raises spreads by 32 basis points (bp), on average. Investors need to be heavily compensated for facing higher credit risk. Moving on to bond-specific features, we find that duration, age and coupon are all positively related to spreads, in line with theory. One more year of interest rate risk entails a 5 bp higher spread; a bond one year older (less "on-the-run") delivers a yield of half a basis point higher; and, having a coupon of one more 1 pp increases corporate yields by 10 bp due to higher income. In turn, the amount outstanding (a measure of size and therefore supply) has a negative effect; this makes sense, as we would expect bonds with a greater supply to bear lower yields.

Next, we turn to our computed variables. We start with equity return volatility, a measure of firm value. In line with our prior, it is positive at high levels of statistical significance. An increment of 1 pp in the standard deviation of the company's stock return pushes spreads upwards by approximately 18 bp, in line with the original work cited in Section 3.2. Our second variable is the Roll measure of liquidity which also performs well, bearing a positive sign at the 1% level. The coefficient predicts a rise of 0.43 bp in corporate spreads for every basis point increase in our *synthetic* bid-ask spread. The order of magnitude of our result is strikingly similar to that of the *observed* bid-ask spread in the work of Chen et al (2007). They find a coefficient of about 0.42 bp.

Having validated that all controls and historical determinants of corporate spreads behave as expected, we now focus on our hypothesis regarding carbon transition risk. Specification (2) shows that, at the 5% level, scope 1 emissions predict corporate spreads – evidence of a *credit risk-adjusted* carbon premium. Concretely, a 1% increase in GHG directly emitted entails a 0.02 basis point yield increment. We can also look at a positive message: what happens when firms reduce their emissions? For instance, by halving their direct carbon emissions (ie reducing them by 50%), firms can reduce their funding costs by 1.1 basis points, on average. Despite statistical significance, the economic effect seems low.

Effect of carbon emissions on corporate spreads

Graph 3.2

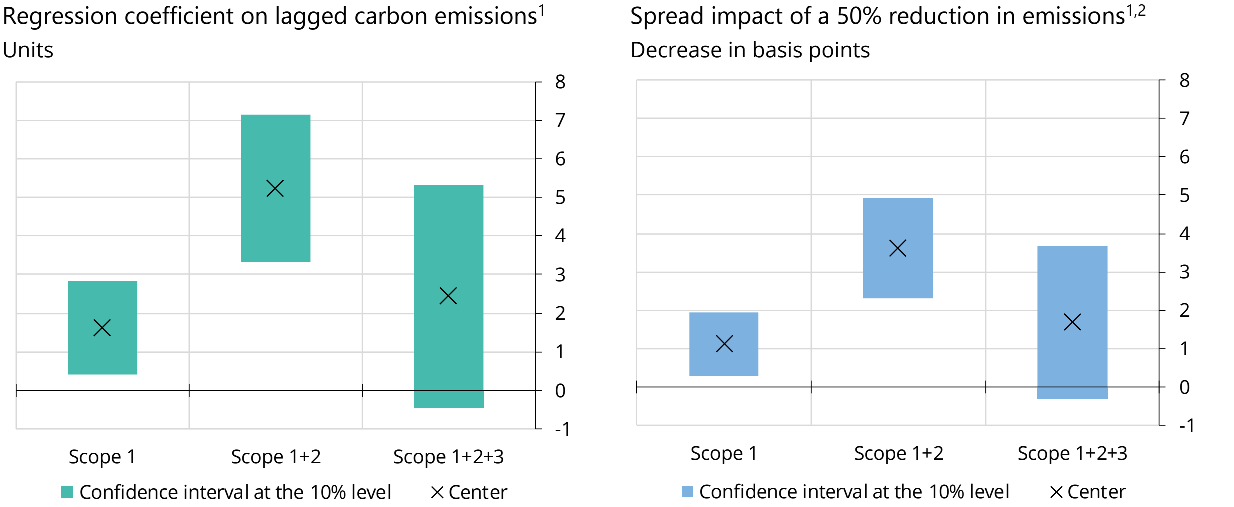

[1] If the bar touches zero, the null hypothesis that the coefficient is zero cannot be rejected. [2] Computed as the coefficient $\hat{\beta}_{P,Carbon}$ on carbon emissions, multiplied by $ln(0.5)$.

Sources: Bloomberg; Refinitiv; Trucost; authors' calculations.

Turning to specification (3), which puts direct *and* indirect emissions together, we see both the statistical power and the economic significance rise. A joint 50% reduction of scope 1 and 2 emissions predicts a 3.6 basis points decrease in the cost of debt at the 1% confidence level. This result makes our carbon premium findings more meaningful. Finally, we note that specification (4), which covers indirect emissions along the value chain, bears a lower coefficient and strips out any statistical significance. We take this result as confirmation of our word of caution about using scope 3 emissions – which are neither widely available nor consistent across providers – as an explanatory variable.

Non-energy-intensive vs energy-intensive carbon premium

Table 3.2

	(1)	(2)	(3)	(4)	(5)	(6)
	Scope 1 emissions			Scope 1+2 emissions		
ln(emissions)	0.86	3.64*		4.54***	4.40*	
	[0.74]	[2.13]		[1.16]	[2.58]	
Non-energy- intensive x ln(emissions)			0.38			3.88***
			[0.73]			[1.17]
Energy- intensive x ln(emissions)			6.96***			8.94***
			[2.39]			[2.93]
Default probability (%)	29.24***	38.25***	31.89***	29.26***	38.22***	31.87***
	[1.25]	[3.37]	[1.26]	[1.26]	[3.37]	[1.26]
Duration	5.10***	6.53***	5.22***	5.11***	6.53***	5.23***
	[0.12]	[0.28]	[0.12]	[0.12]	[0.28]	[0.12]
Age	0.39***	1.28***	0.56***	0.39***	1.28***	0.56***
	[0.15]	[0.31]	[0.14]	[0.15]	[0.31]	[0.14]
Coupon	10.87***	9.36***	10.52***	10.87***	9.35***	10.52***
	[0.52]	[0.99]	[0.47]	[0.52]	[0.99]	[0.47]
ln(amount outstanding)	-2.26***	-6.63***	-2.91***	-2.27***	-6.64***	-2.92***
	[0.31]	[1.28]	[0.32]	[0.31]	[1.28]	[0.32]
Equity return volatility	15.57***	28.75***	17.68***	15.63***	28.74***	17.70***
	[0.98]	[2.63]	[0.93]	[0.98]	[2.62]	[0.93]
Liquidity	0.40***	0.47***	0.43***	0.40***	0.47***	0.43***
	[0.02]	[0.04]	[0.02]	[0.02]	[0.03]	[0.02]
Callable	-6.82***	-12.88***	-7.97***	-6.80***	-12.88***	-7.94***
	[1.04]	[2.75]	[0.99]	[1.04]	[2.75]	[0.99]
Number of bonds	6330	1269	7599	6330	1269	7599
Observations	221,667	42,015	263,682	221,782	42,015	263,768

*** p<0.01, ** p<0.05, * p<0.1. Standard errors in brackets; clustered at the security level.

Sources: Bloomberg; Refinitiv; Trucost; authors' calculations.

Our results are summarised visually in Graph 3.2. The credit risk-adjusted premium appears statistically significant (left-hand panel, our regression coefficients) when we account for direct and indirect emissions that are required disclosures by the Greenhouse Gas Protocol. And the economic effect (right-hand panel, basis points) appears highest when the sum of scope 1 and 2 emissions are considered.

3.3 Exploring sector effects

In Section 2, we showed an important step difference in the order of magnitude of emissions between companies considered "energy-intensive" and those that are not. So, should bonds from all firms bear the same carbon premium? In this subsection, we seek to answer this question.

To analyse differences in the carbon premium between energy-intensive and non-energy-intensive sectors, we conduct two different exercises:

a. We split the bond sample into two subsamples. One with securities from non-energy-intensive firms and another with bonds from energy-intensive ones. This allows us to vary the coefficients on the control variables depending on the firm's sector.

b. We run a full-sample exercise which interacts our carbon emissions variable with an energy-intensive sector dummy. The value of the dummy is equal to 1 when the company belongs in the category and zero otherwise.

Table 3.2 shows the results for exercises (a) and (b) across the different groups of emissions. We start by describing the subsample results for scope 1. Specifications (1) and (2) bear very different coefficients in front of the log-emissions variable. For non-energy-intensive sectors (model 1), the coefficient is below one and does not appear statistically significant; yet for energy-intensive sectors (model 2), $\hat{\beta}_{P,Carbon}$ grows fourfold and becomes significant at the 10% level. When we apply the dummy variable (model 3), we find a similar result: the carbon premium is much higher for firms in energy-intensive sectors.

This result changes somewhat when looking at scope 1 and 2 emissions jointly. Models 4 to 6 show how including indirect emissions in the computation gives its coefficient statistical significance across all kinds of firms, regardless of their energy consumption. Though the subsample results show similar premia for both cases (models 4 and 5), our dummy specification in particular (model 6) offers a carbon premium at least twice as big for bonds from energy-intensive companies. Finally, by looking at models 7 to 9, we once again see that using scope 3 emissions in our modelling introduces awkward dynamics (eg a negative sign on $\hat{\beta}_{P,Carbon}$ for non-energy-intensive bonds (model 7)).

The risk-adjusted carbon premium by sector

Graph 3.3

Effect of carbon emissions on corporate spreads[1]
Coefficient

Spread decrease implied by −50% in emissions
Basis points

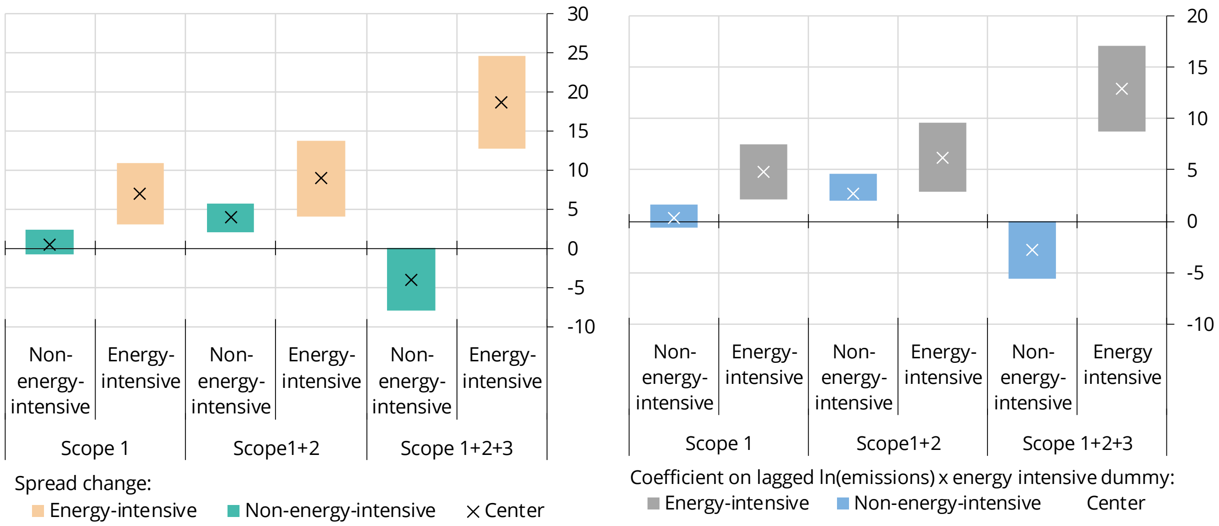

[1] Specifications where lagged log-carbon emissions are interacted with an indicator variable equal to 1 when the firm is from an energy-intensive sector. Bars denote a confidence interval at the 10% level. If the bar touches zero, the null hypothesis that the coefficient is zero cannot be rejected. [2] Computed as the coefficient $\hat{\beta}_{P,Carbon}$ on carbon emissions, multiplied by $ln(0.5)$.

Sources: Bloomberg; Refinitiv; Trucost; authors' calculations.

A graphical summary of specifications (b) with the interaction is offered in Graph 3.3. We highlight a few differences with our previous exercise: first, there are important changes in the carbon premium between companies considered energy-intensive and those that are not. The coefficients (left-hand panel) are at least twice the size for more polluting ("browner") firms.

As a consequence, the spread impact is much increased. A reduction of 50% in 1+2 emissions could help energy-intensive firms' bonds trade 9 bp cheaper. This is equivalent to a rating upgrade of 0.8 notches.[36] Given the relevance of both scope 1 and scope 2 emissions in compulsory reporting, we consider model (6) the key finding of this section. Zooming in on firms in energy-intensive sectors reveals a more meaningful preference channel than found in the overall sample. The impact within the set of energy-intensive firms is statistically significant and of non-negligible economic importance.

[36] This result is the average of individual notch changes across all bonds. To estimate each individual notch change, the 9 bp impact from a 50% change in emissions is divided by the differential between two (mean) spreads: that of the bond's credit rating and that of the adjacent notch below. Our sample comprises bonds with ratings AAA– to C on the Fitch scale.

3.4 The effects of maturity

One defining characteristic of bond spreads is the existence of the term structure. As Graph 3.1 showed, there is a relationship between corporate yields and bond maturity. In this section, we ask ourselves whether carbon risk compensation may also be related to the term of each instrument. The analysis is possible thanks to the security-level approach we have taken in our models.

Looking by maturity

Table 3.3

	(1) Scope 1 emissions	(2) Scope 1+2 emissions	(3) Scope 1+2+3 emissions
Maturity < 5 years x ln(emissions)	0.11	3.87***	1.16
	[0.72]	[1.14]	[1.71]
Maturity 5-10 years x ln(emissions)	2.24***	5.99***	3.14*
	[0.71]	[1.13]	[1.70]
Maturity 10-15 years x ln(emissions)	3.03***	6.82***	3.95**
	[0.72]	[1.13]	[1.70]
Maturity 15-20 years x ln(emissions)	4.05***	7.87***	4.94***
	[0.73]	[1.14]	[1.70]
Maturity 20-25 years x ln(emissions)	3.64***	7.58***	4.70***
	[0.72]	[1.13]	[1.70]
Maturity 25-30 years x ln(emissions)	2.75***	6.83***	4.07**
	[0.73]	[1.14]	[1.70]
Maturity > 30 years x ln(emissions)	3.05***	7.16***	4.41***
	[0.76]	[1.15]	[1.71]
Number of bonds	7599	7599	7599
Observations	263,682	263,797	263,797
R-squared	0.84	0.84	0.84

*** p<0.01, ** p<0.05, * p<0.1. Standard errors in brackets; clustered at the security level. Coefficients on other variables omitted for brevity.

Sources: Bloomberg; Refinitiv; Trucost; authors' calculations.

To answer our question, we classify bonds according to maturity. To do so, we create a dummy variable within buckets, using five-year steps. We start with bonds of less than five years in maturity, then with bonds between five and 10 years, and so forth – our last bucket comprises all bonds above 30 years. Next, we rerun our model, by interacting carbon emissions with each of these maturity buckets.

Table 3.3 presents the results, which provide evidence of a term structure of carbon premia. Across models 1, 2 and 3 – which differ in emission scopes being taken into account – we find statistical significance in most coefficients on the interaction. In other words, maturity and emissions *together* help explain the cross-section of corporate bond spreads. And the magnitude varies by term. Looking at the magnitudes, carbon risk appears to impact bonds in the 15–20 years bucket the most and shorter maturities (securities with <5 years) the least (see statistical significance).

To better showcase our results, Graph 3.4 contains the decreases in spread implied by the halving of firm-level emissions under these models. By controlling scope 1 GHG (left-hand panel), firms may reduce their financing costs by some 0–3 bp, depending on maturity, on average. The effect is up to 5.5 bp – in expectation –

for scope 1 and 2 emissions together (centre panel). However, the confidence intervals denoted by the bars show that an effect of up nearly 7 bp for the belly of the curve is possible. With slightly more uncertainty (right-hand panel, wider blue bars) this effect is up to 4.7 bp, on average for maturities between 15 to 20 years when all emission scopes are taken into account.

Term structure of credit risk-adjusted carbon premia[1]

Spread decrease induced by a 50% reduction in carbon emissions, in basis points Graph 3.4

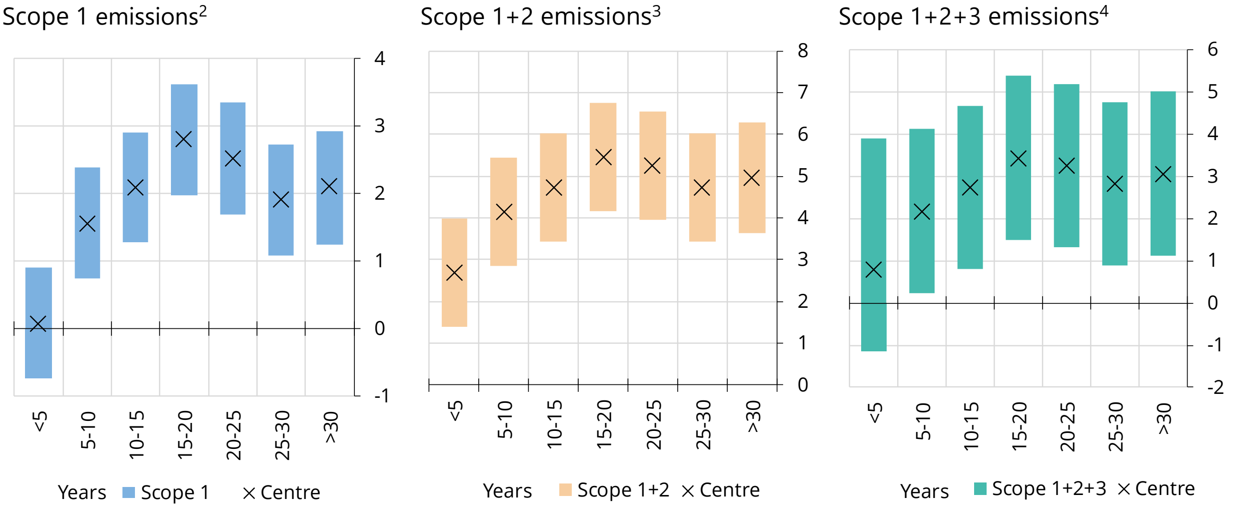

[1] Computed as the coefficient $\hat{\beta}_{P,Carbon}$ on carbon emissions, multiplied by $ln(0.5)$.

[2] Corresponds to model (1) in Table 3.3. [3] Corresponds to model (2) in Table 3.3. [4] Corresponds to model (3) in Table 3.3.

Sources: Bloomberg; Refinitiv; Trucost; authors' calculations.

Up to this point, we have found that both sector and maturity matter, independently. What happens when we explore both effects simultaneously? In what follows, we interact both indicator variables with carbon emissions, to investigate whether there are two term structures: one for energy-intensive firms and another for their complement. This time, we skip the formalities, going straight to our term structure computations for the regulatory emissions (scope 1 and scopes 1+2). We show these in Graph 3.5.

Our key finding here is that, indeed, cross-industry results do mask important differences in the term structures of non-energy-intensive (Graph 3.5, left-hand side) and energy-intensive firms (right-hand side). Focusing on scope 1+2 emissions (two bottom panels) we see that firms' bonds trade up to 4.2 basis points lower when total emissions are cut back 50% (Graph 3.5, bottom left-hand panel, white crosses, maximum value). The result is strikingly higher for energy-intensive firms (Graph 3.5, bottom right-hand panel, white crosses), where the effect can be up to 8.2 bp. In fact, our confidence intervals take our estimates – which are all statistically significant at the 1% level – to a spread effect of up to 13 bp. As in the overall sample results, the effects are more pronounced for the maturities after 15 years, the 15–20 year bucket being the greatest. The aspects of this so-called term structure of carbon premia appears rather hump-shaped.

A second glance at the term structure: by sector[1,2]

Spread decrease induced by a 50% reduction in emissions, in basis points | Graph 3.5

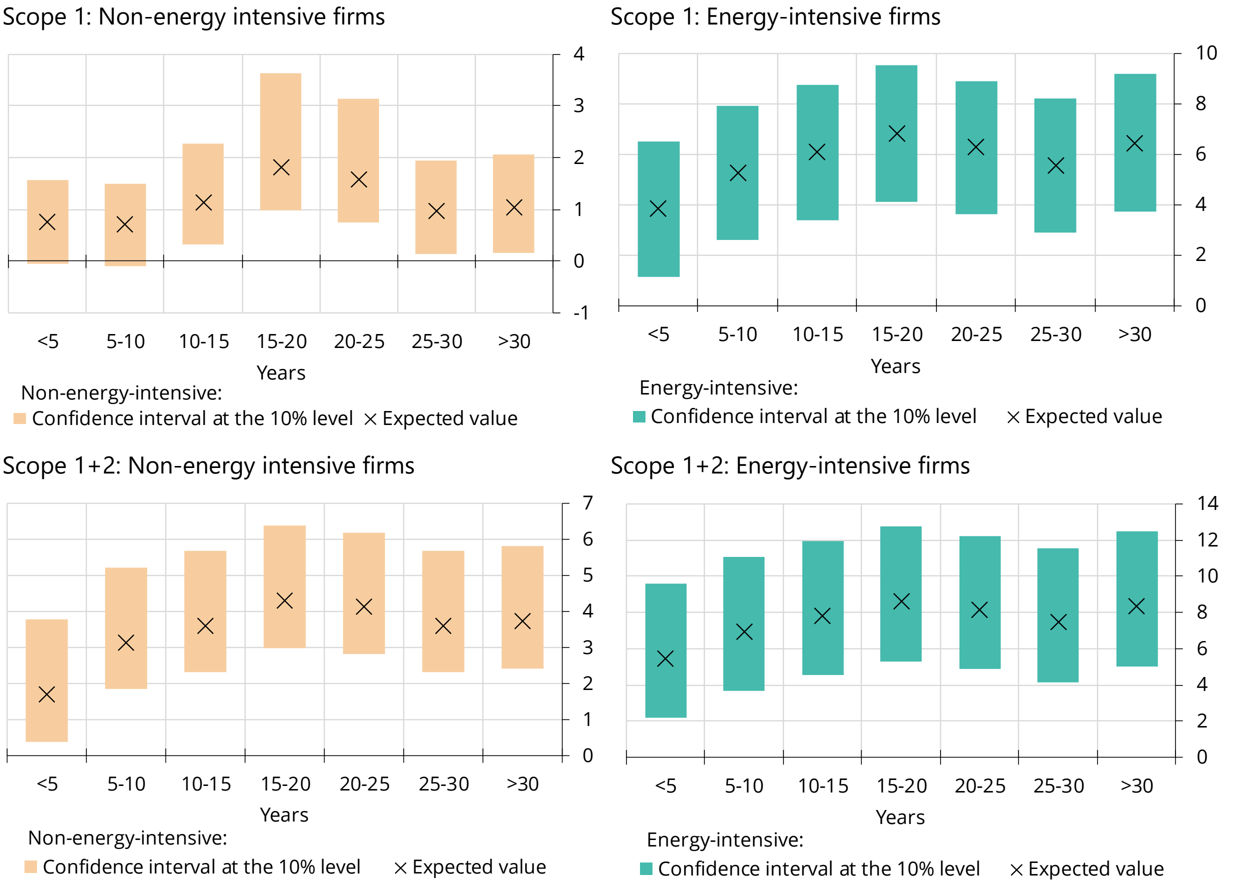

[1] The impact is calculated as the product between the coefficient on $ln(\text{emissions})$ and a 50% change in emissions (ie as $\hat{\beta}_{P,Carbon} \times ln(0.5)$). [2] The results corresponds to model (2) in Table 3.4.

Sources: Bloomberg; Refinitiv; Trucost; authors' calculations.

Table 3.4 presents a regression table with all specifications, grouping results by energy-intensiveness, where model (2) corresponds to Graph 3.5 above. As usual, results for scope 1 on its own (model (1)) show a smaller carbon premia than those for scope 1 and 2 together. The result holds regardless of whether a company is energy-intensive or not. When we look at the coefficients which include scope 3 GHG, the term structure appears to be inverted for non-energy-intensive firms (ie decreasing with maturity) and defined in the [2,4] bp range. For the energy-intensive case the function is mostly upward sloping and defined roughly in the range [12,16]. However, there are large differences in statistical significance for this model. For the majority of maturity groups, the coefficient fails to pass any significance test when the bond spread is from a non-energy-intensive company (upper half of the table, column 3), with some values bearing a negative sign. We interpret this erratic behaviour as one more piece of evidence that using scope 3 emissions for empirical analysis requires careful deliberation.

Twin term structures under the preference channel

Table 3.4

	(1) Scope 1 emissions	(2) Scope 1+2 emissions	(3) Scope 1+2+3 emissions
Non-energy-intensive x ln(emissions) x			
Maturity < 5 years	-1.09	2.45**	-5.76***
	[0.71]	[1.14]	[1.83]
Maturity 5-10 years	1.02	4.51***	-3.69**
	[0.71]	[1.13]	[1.83]
Maturity 10-15 years	1.64***	5.19***	-2.93
	[0.72]	[1.13]	[1.83]
Maturity 15-20 years	2.61***	6.19***	-1.93
	[0.72]	[1.14]	[1.83]
Maturity 20-25 years	2.26***	5.94***	-2.14***
	[0.73]	[1.14]	[1.82]
Maturity 25-30 years	1.40**	5.22***	-2.73
	[0.73]	[1.14]	[1.83]
Maturity > 30 years	1.49**	5.37***	-2.57
	[0.76]	[1.15]	[1.84]
Energy-intensive x ln(emissions) x			
Maturity < 5 years	5.52**	7.90***	18.33***
	[2.35]	[2.90]	[3.59]
Maturity 5-10 years	7.56***	10.03***	20.07***
	[2.34]	[2.89]	[3.58]
Maturity 10-15 years	8.74***	11.30***	21.03***
	[2.34]	[2.89]	[3.58]
Maturity 15-20 years	9.82***	12.45***	22.02***
	[2.36]	[2.91]	[3.59]
Maturity 20-25 years	9.03***	11.73***	21.69***
	[2.32]	[2.88]	[3.58]
Maturity 25-30 years	8.01***	10.77***	20.94***
	[2.34]	[2.89]	[3.58]
Maturity > 30 years	9.28***	12.04***	22.11***
	[2.40]	[2.93]	[3.60]
Number of bonds	7599	7599	7599
Observations	263,682	263,797	263,797
R-squared	0.84	0.84	0.84

*** p<0.01, ** p<0.05, * p<0.1. Standard errors in brackets; clustered at the security level. Coefficients on other variables omitted for brevity.

Sources: Bloomberg; Refinitiv; Trucost; authors' calculations.

Our key results of this section are models (1) and (2) of Table 3.4 and encapsulate our novel finding of a term structure or "curve" of credit risk-adjusted carbon premia, which is hump-shaped and depends on a given firm's sector.[37] While the jury is still out on the curve's shape, we offer two conjectures:

a) The first is the long-term nature of environmental risks, which, despite requiring critical action today, will become inevitable in only a few years.

[37] It is "credit risk-adjusted" because, in our analysis of the preference channel, we control for the probability of defaults.

Indeed, as the race to net zero continues, it is likely that companies that pollute the most will be the first to face dramatic investor dispreference (eg fire-sale risks) should transition risks which are seemingly far away appear more material to the public.[38]

b) The second is preferred habitat. The underlying assumption is that demand and supply forces play different roles across different sectors of the curve. It may well be that investors trading on information about the environmental impact of a company may operate more in particular maturities over others. In the light of our results, this may not be the very short term. For example, pension funds, which are increasingly aware of sustainable investing, tend to have a preference for longer-term bonds in order to match their liabilities; but they may not opt for the ultra-long segments due to liquidity and interest rate risk concerns.

The differentiated impact across industries probably reflects investors' greater scrutiny of firms viewed as brown – emitting much more significant amounts of GHGs into the atmosphere is highly penalised by the market. Put together, our findings support the presence of a so-called preference channel for carbon risk in the corporate bond market.

4. The risk channel

In this section, we explore how a firm's carbon footprint affects bond spreads through the credit risk channel. In other words, we test whether a firm's default probability – established as a key determinant of corporate bond spreads in Section 3 – reflects any exposure to transition risk as measured by carbon emissions. Our hypothesis is that firms with higher GHG emissions are more exposed to transition risk and are therefore more likely to default, all else equal. Our conjecture is in line with practices at banks and rating agencies who factor a firm's environmental impact into their credit risk assessments.

4.1 The model

To assess the impact of emissions on the probability of a given firm defaulting, we run the following panel regression:

$$\tilde{P}_{i,t} = \beta_{R,carbon} \times \ln(\text{Emissions}_{i,t-12}) + \delta' X_{i,t} + FE + \varepsilon_{i,t}. \quad (4)$$

where the left-hand variable is the five-year-ahead annualised probability of default for firm i at time t, gathered from Bloomberg. This probability differs from $P_{i,t}$ in equation (3) in that $\tilde{P}_{i,t}$ is its logit transformation or "log-odds of default", which we compute as:

$$\tilde{P}_{i,t} = \ln\left(\frac{P_{i,t}}{1 - P_{i,t}}\right) \quad (5)$$

38 This particular result should be differentiated from that with a credit risk interpretation. The effect of carbon emissions on a firm's perception of default is explored in the following section, which covers the (credit) risk channel.

We apply this transformation so that our regressand is not bounded between 0 and 1. This way, our specification is also consistent with the commonly used logit models capturing default events (eg Duffie et al (2007)). In equation (4), $X_{i,t}$ is a set of firm-level control variables, including: size of assets, long-term debt-to-assets ratio, earnings-to-assets ratio, capital-to-assets ratio and return-on-assets (RoA). FE represents a vector of time and sector fixed effects and $\varepsilon_{i,t}$ is the residual. In this new specification, $\beta_{R,Carbon}$, the coefficient in front of the one-year lagged carbon emissions, corresponds to our estimate of the risk channel. And, if our hypothesis holds, $\beta_{R,Carbon}$ should be positive.

Similar to our analysis in the previous section, we estimate the model with monthly data starting in January 2017 and consider direct (scope 1) emissions and indirect emissions jointly (scopes 1+2 and scopes 1+2+3).

4.2 Full sample results

We first run the regression using the full sample and show our results in Table 4.1. As in Section 3, we offer four baseline specifications. Columns (2) to (4) correspond to the models capturing different carbon emission scopes. Column (1) is a benchmark model that leaves out carbon emissions but is otherwise identical to equation (4) above.

Effects of carbon emissions on firm PD

Table 4.1

	(1)	(2) Scope 1 emissions	(3) Scope 1+2 emissions	(4) Scope 1+2+3 emissions
ln(emissions)		0.06***	0.07***	0.06***
		[0.01]	[0.01]	[0.01]
ln(assets)	-0.14***	-0.19***	-0.21***	-0.20***
	[0.01]	[0.01]	[0.01]	[0.01]
Long-term debt/assets	1.53***	1.57***	1.56***	1.58***
	[0.07]	[0.07]	[0.07]	[0.07]
Earnings/assets	-0.13***	-0.13***	-0.13***	-0.13***
	[0.02]	[0.02]	[0.02]	[0.02]
Capital/assets	-0.43***	-0.49***	-0.53***	-0.51***
	[0.04]	[0.04]	[0.04]	[0.04]
Return on assets (%)	-0.04***	-0.04***	-0.04***	-0.04***
	[0.001]	[0.002]	[0.002]	[0.002]
Number of firms	2910	2831	2831	2831
Observations	150,176	140,774	140,858	140,858
R-squared	0.48	0.5	0.5	0.5

*** p<0.01, ** p<0.05, * p<0.1. Standard errors in brackets; clustered at the firm level.

Sources: Bloomberg; Refinitiv; Trucost; authors' calculations.

The benchmark model yields a lower R^2 than models with carbon emissions, suggesting that firm-level greenhouse gases do indeed play a role in credit risk assessments. For all four models, the coefficients in front of the control variables bear the right signs and are statistically significant to the 1% level. For instance, as expected, a broader balance sheet, higher earnings, stronger capital and increased RoA are all related to lower probabilities of default. The opposite is true for the ratio of long-term debt to assets. Moreover, the magnitude of these coefficients is broadly similar across the different specifications, implying that our measure of carbon emissions represents information not contained in the set of other firm-specific characteristics considered.

Zooming in on models that consider carbon emissions, these baseline results confirm our hypothesis – a firm's emitted level of CO_2 and other greenhouse gases has an adverse impact on its probability of default. For all carbon emission measures, $\beta_{R,carbon}$ is estimated to be positive and statistically significant with at least 99% confidence. The magnitude of our estimate $\hat{\beta}_{R,carbon}$ is roughly similar regardless of which scopes are included in the computation. This finding suggests there may be limited differentiation between direct and indirect carbon emissions when it comes to assessing transition risk.

Given that our model is one of log-odds of default, how should these coefficients be interpreted in terms of default probability levels? Take scope 1 carbon emissions, for example. Halving emissions would translate to a 0.03 decrease in log-odds $\tilde{P}_{i,t}$. However, since the relationship between $\tilde{P}_{i,t}$ and $P_{i,t}$ is non-linear, the impact of carbon emissions on default probability (and thus, option-adjusted spreads) is not a constant. To compute the resulting effect on probability levels, we instead proceed as below.

First, we estimate the change on log-odds $\Delta\tilde{P}_i$ induced by a change in log-emissions ΔE. This is obviously a function of our coefficient $\hat{\beta}_{R,carbon}$:

$$\Delta\tilde{P}_i = \hat{\beta}_{R,carbon} \times \Delta E \tag{6}$$

Next, and for each probability level p_i, we compute its log-odds and then add the change in log-odds $\Delta\tilde{P}_i$ to obtain a new log-odds probability level $\tilde{P}_i'$. From equation (5), this is:

$$\tilde{P}_i' = \Delta\tilde{P}_i + \ln\left(\frac{p_i}{1-p_i}\right) \tag{7}$$

Finally, we apply the inverse logit transformation to our estimate $\tilde{P}_i'$ to derive the resulting default probability level p_i'. We lastly compare this with our original probability level p_i to get the estimated change in probability level Δp_i:

$$\Delta p_i = p_i' - p_i = \frac{1}{1+e^{\tilde{P}_i'}} - p_i \tag{8}$$

For our purposes, we define p_i continuously in the range [0, 0.32]. In other words, we only consider probabilities of default of up to 32% as it is the upper bound for this variable in our sample (see Section 2.1 for details).

The left-hand panel of Graph 4.1 plots our resulting estimates for each level of default probability when emissions are halved. Within the observed range, we notice that the impact increases with default likelihood. At the maximum probability of 32%, halving emissions – both direct and combined – translates to a 1 pp decrease in probability. These estimates can be used to derive the impact on option-adjusted spreads, which we will discuss in the next section.

Decline in five-year default probabilities induced by a 50% reduction in emissions

In percentage points — Graph 4.1

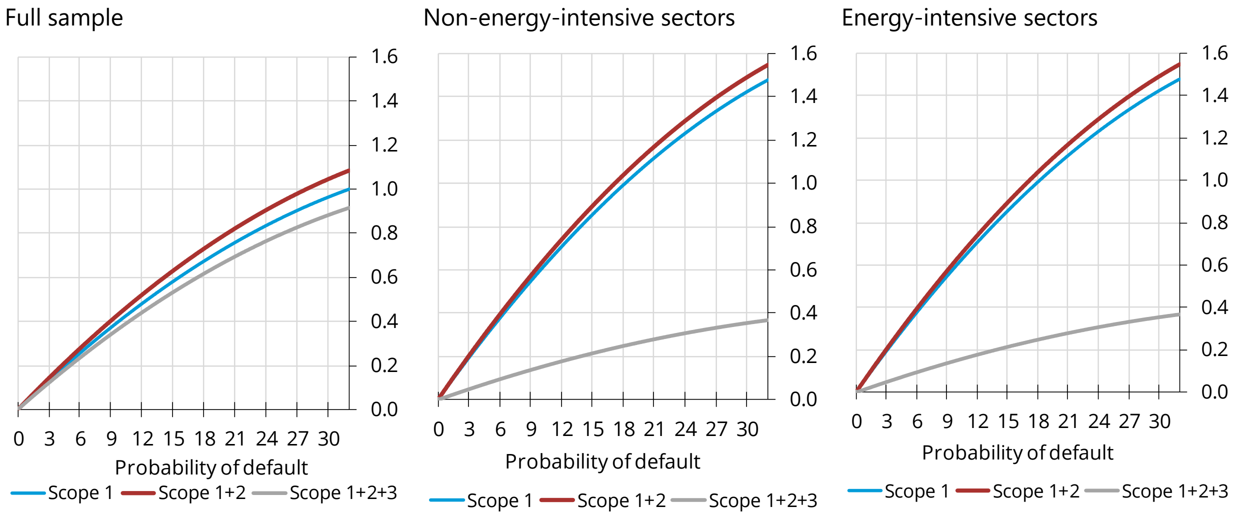

[1] The effect of a 50% decrease in carbon emissions is computed continuously for each probability level on the scale 0 to 32%.

Sources: Bloomberg; Refinitiv; Trucost; S&P Capital IQ; authors' calculations.

4.3 Energy-intensive vs non-energy-intensive sectors

As in the preference channel, we also wish to test whether the impact of emissions on corporate default is only viable in a few sectors, specifically the ones commonly viewed as "brown" industries with heavy GHG emissions. To this end, we consider two analyses, as before:

a. A subsample analysis, where we run regressions exclusively for firms belonging to energy-intensive sectors and for those in non-energy-intensive sectors, respectively.

b. A model which adds an interaction term between carbon emissions and whether a firm is from an energy-intensive sector or not. This adds a dummy to the regressor list in equation (4) above. Our classification of energy-intensive and non-energy-intensive sectors is identical to the previous section.

Table 4.2 shows the results. Columns (1), (4) and (7) are estimates based on the subsample of non-energy-intensive sectors; columns (2), (5) and (8) are estimates using the subsample of energy-intensive sectors; and columns (3), (6) and (9) are derived from models with the dummy interaction term.

Our results suggest that the impact of carbon emissions through the risk channel prevails across different sectors with a greater effect on energy-intensive sectors. In both subsamples, carbon emissions play a negative and statistically significant role in the probability of default. The results from regressions with interactions are consistent with the results from subsample analysis. For scope 1 and scope 1+2 emissions, their impact on the probability of default is larger in energy-intensive sectors. For scope 1 emissions, a 50% increase of the emissions would translate to a 0.04 increase in $\tilde{P}_{i,t}$ in non-energy-intensive sectors and a 0.07 increase in $\tilde{P}_{i,t}$ in energy-intensive sectors. These correspond to 0.9% and 1.5% increases in the probability of default at probability of 32% (centre and right-hand panels of Graph 4.1). The impact of scope 1+2 emissions is quite similar to that of scope 1 emissions: 1% and 1.6% respectively for non-energy–intensive and energy–intensive sectors. Interestingly, for scope 1+2+3 emissions, the risk channel impact on non-energy–intensive sectors is similar to that of other emission measures but the impact on energy-intensive sectors is relatively smaller. This could partly reflect data quality issues with scope 3 emissions.

Table 4.2

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Scope 1 emissions			Scope 1+2 emissions			Scope 1+2+3 emissions		
ln(emissions)	0.06***	0.09***		0.06***	0.10***		0.07***	0.02***	
	[0.01]	[0.02]		[0.01]	[0.03]		[0.01]	[0.03]	
Non-energy-intensive x ln(ems.)			0.06***			0.06***			0.06***
			[0.01]			[0.01]			[0.01]
Energy-intensive x ln(emissions)			0.08***			0.09			0.05**
			[0.01]			[0.01]			[0.02]
ln(assets)	-0.19***	-0.24***	-0.19***	-0.20***	-0.25***	-0.21***	-0.21***	-0.17***	-0.20***
	[0.01]	[0.03]	[0.01]	[0.01]	[0.03]	[0.01]	[0.01]	[0.04]	[0.01]
Long-term debt/assets	1.57***	1.63***	1.57***	1.56***	1.64***	1.56***	1.55***	1.85***	1.58***
	[0.08]	[0.21]	[0.07]	[0.08]	[0.21]	[0.07]	[0.08]	[0.26]	[0.07]
Earnings/assets	-0.12***	-0.19***	-0.13***	-0.11***	-0.18***	-0.13***	-0.12***	-0.13***	-0.13***
	[0.02]	[0.07]	[0.02]	[0.02]	[0.07]	[0.02]	[0.02]	[0.08]	[0.02]
Capital/assets	-0.49***	-0.48*	-0.49***	-0.53***	-0.51*	-0.53***	-0.52***	-0.39***	-0.51***
	[0.04]	[0.27]	[0.04]	[0.04]	[0.27]	[0.04]	[0.04]	[0.25]	[0.04]
Return on assets (%)	-0.04***	-0.02***	-0.04***	-0.04***	-0.02***	-0.04***	-0.04***	-0.03***	-0.04***
	[0.002]	[0.003]	[0.002]	[0.002]	[0.003]	[0.002]	[0.002]	[0.004]	[0.002]
Number of firms	2,443	388	2,831	2,443	388	2,831	2,466	365	2,831
Observations	121,124	19,650	140,774	121,208	19,650	140,858	122,840	18,018	140,858
R-squared	0.49	0.54	0.5	0.49	0.54	0.5	0.48	0.56	0.5

*** p<0.01, ** p<0.05, * p<0.1. Standard errors in brackets; clustered at the firm level.

Sources: Bloomberg; Refinitiv; Trucost; authors' calculations.

5. Combining the two channels

In our previous sections, we derived estimates for two types of carbon premium in corporate bonds: the credit risk-adjusted carbon premia and the credit risk carbon premia.

This section explores the *total* carbon premium that comes out of combining the effects above, which represent the preference and risk channels. To compute total premia, we first need to translate the impact of emissions on default probabilities (the risk channel) into an effect on option-adjusted spreads. This can be achieved by multiplying the estimated risk channel effect on default probabilities (Δp_i from equation (8)) by the effect of this probability on spreads ($\hat{\beta}_P$ from equation (3)).

Graph 5.1 showcases an example. For the full sample (left-hand panel) and for non-energy-intensive sectors (centre panel), a halving of carbon emissions would narrow corporate spreads by as much as 30 basis points for both direct and combined emissions (several scopes together). For energy-intensive firms, the impact is much larger. For scope 1 and scopes 1+2, a halving of carbon emissions could translate to a near 60 basis points decline in spreads.

Decline in OAS induced by a 50% reduction in carbon emissions

In basis points

Graph 5.1

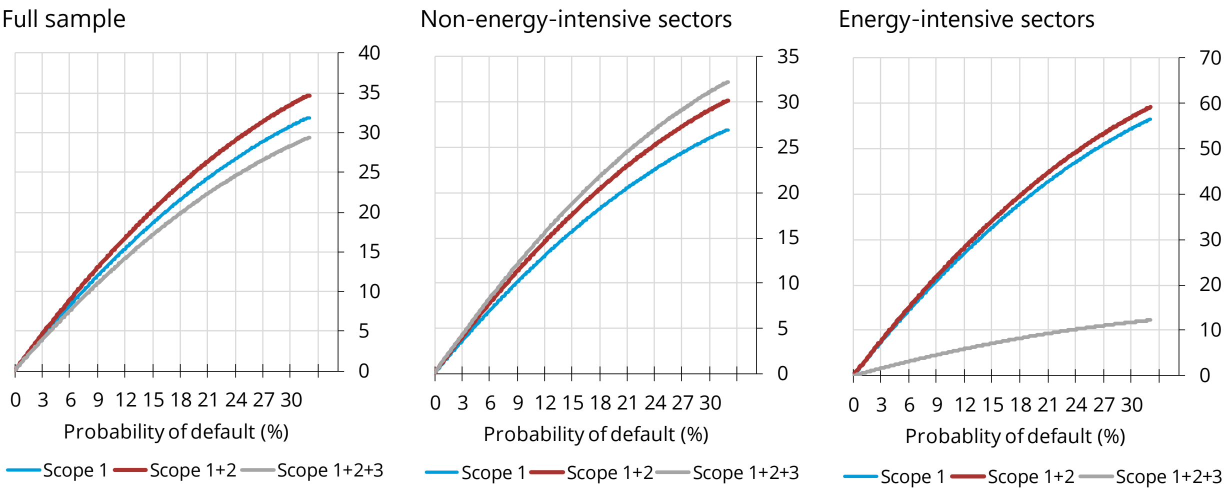

Sources: Bloomberg; Refinitiv; Trucost; S&P Capital IQ; authors' calculations.

With the effect of the risk channel expressed in terms of the spread, we can now combine it with the impact of the preference channel. To this end, we consider a typical firm, whose probability of default equals the sample average, which is 0.56%. Based on this, Graph 5.2 plots the total impact on spreads of a halving in firm-level GHG across different emission measures and sectors. For the full sample (first bar of each graph section), the total impact ranges from 2 to 4.5 basis points, depending on the emission scopes considered. Looking at the bar colours, we note that total premia are mainly explained by the preference channel (green bars) rather than the risk channel (blue bars). Of course, this is a function of the probability of default *level* of the company in question. For the average firm, the PD is low (recall it is approximately 0.56%). This attribution of the total premia changes when we look at a different PD level. For instance, when the PD is one standard deviation above the average (1.54%), the contributions from the two channels are on more equal footing.

Comparing across different sectors, the total impact appears larger for energy-intensive firms. Concretely: for a typical firm in an energy-intensive sector, the impact is around 8 basis points for scope 1 and scope 1+2 emissions and more than 13 basis points for scope 1+2+3 emissions combined. In comparison, for a typical firm in the non-energy-intensive category, the impact is at most 3 basis points for scope 1 and 2 combined.

Combined impact on spreads of a 50% reduction in carbon emissions

In basis points Graph 5.2

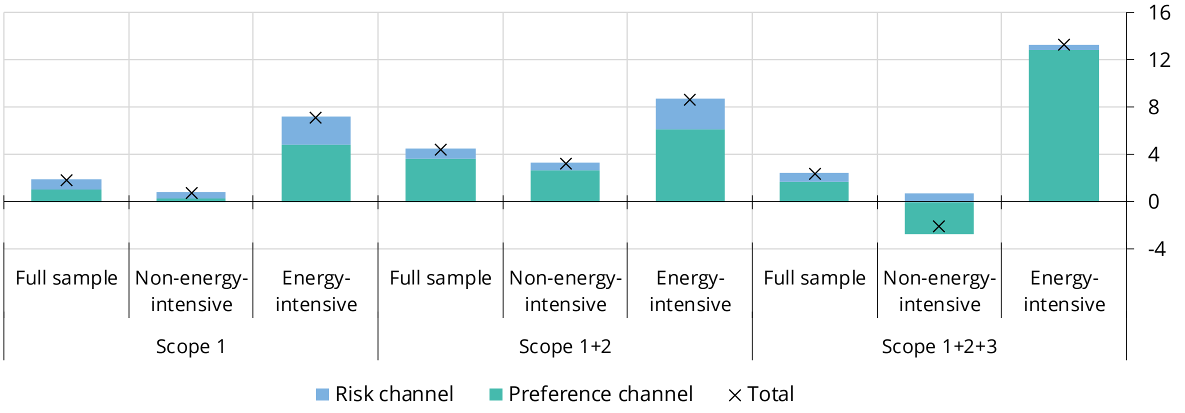

Sources: Bloomberg; Refinitiv; Trucost; S&P Capital IQ; authors' calculations.

We then look into the term structure of total carbon premia. Recall that the term structure coming from the preference channel is hump-shaped. The addition of the risk channel largely preserves this hump shape (results available upon request). This is because the impact of carbon emissions on spreads via credit risk depends on the impact of emissions on a firm's default probability and on the mapping of a firm's default probability to the spreads of bonds issued by the company. The former is independent of maturity, as default probability is gauged at the firm level and not the bond level, while the latter is broadly identical across different maturities.[39] With the upward shift induced by the addition of the risk channel, the total carbon premium in the belly (the maturity bucket being 15–20 years) is in the range of 3.5 to 6 basis points, depending on emission measures.

Last, we separately contrast the term structures of total carbon premia for non-energy-intensive and energy-intensive sectors. Graph 5.3 presents our results for scope 1 (top panels) and scope 1+2 (bottom panels). For brevity, we focus on the results of scope 1+2 emissions – the measure yielding the largest total carbon premia. As can be seen in Graph 5.3 (bottom two panels), the term structure is hump-shaped for both non-energy-intensive and energy-intensive categories. Across the maturity spectrum, total carbon premia are almost twice as large for energy-intensive firms than for non-energy-intensive ones. For bonds maturing in 15–20 years, a halving in scope 1+2 emissions would narrow their spreads by more than 10 basis points. For "greener" firms, the effect is around 5 basis points.

39 We have also estimated $\hat{\beta}_P$ from a more flexible model in which we interact carbon emissions with the maturity bucket indicator. The results are very close. For brevity, they are not presented here but are available upon request.

Term structure of total carbon premia: by sector[1]

Spread decrease induced by a 50% reduction in emissions, in basis points — Graph 5.3

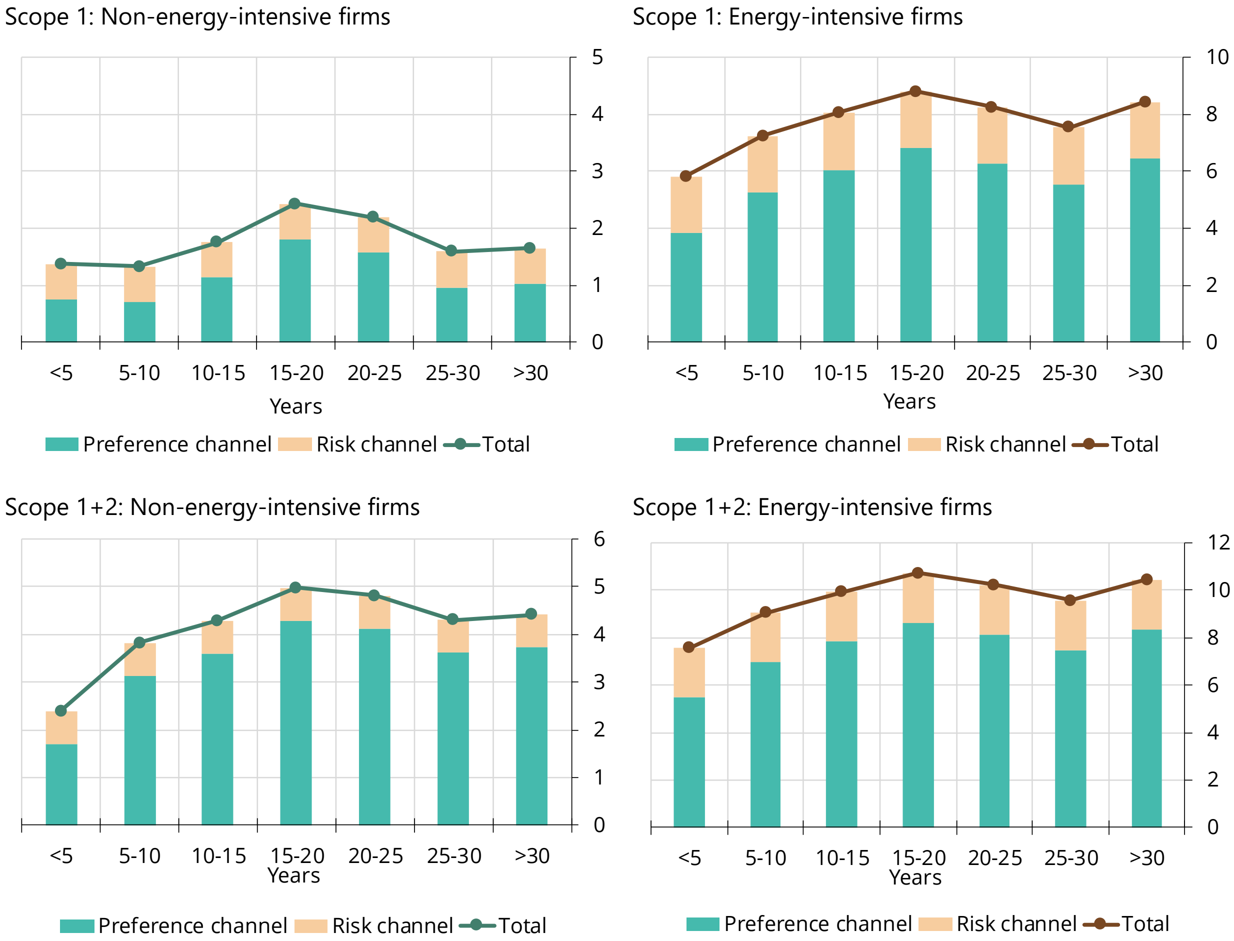

[1] Effect of a 50% reduction in the respective emission scope totals.

Source: Bloomberg; Refinitiv; Trucost; S&P Capital IQ; authors' calculations.

6. Robustness checks

In Sections 3 and 4, we explored whether our findings hold when excluding securities from the most carbon-intensive sectors. This led us to conclude that our results are robust to these formulations and that we can in fact compute two different term structures of total carbon premia. To further validate our findings, we have conducted other robustness checks, which we present in this subsection.

6.1 Robustness checks for the preference channel

We consider two robustness checks for the preference channel, with alternative measures of default probability and liquidity, respectively.

The first check consists of swapping Bloomberg's measure of probability of default for our own computations in the preference channel model. As detailed in Section 2, our own computations are based on the work of Merton (1974).[40] This procedure returns an alternative metric of five-year-ahead default probability, which we use as regressor in place of the Bloomberg ready-made ones.

The regression models using our computed default probabilities are in Table 6.1. Before looking at the carbon emission results, we review any changes to the model without carbon emissions (column 1). What are the most important changes? In particular, a rise of 1 percentage point in the risk-neutral default probability can be translated to an increase of 3 basis points in bid-ask spreads. This is about a 10th of the effect found in physical default probabilities. The result is intuitive, since in the risk-neutral world – and because investors are in aggregate risk-averse – prices imply higher probabilities to negative scenarios than they do to positive ones. Another important change is in equity return volatility, which now has a coefficient 1.7 times the original. The rest of the coefficients show similiar magnitudes across the board. The impact of duration, age, coupons, outstanding amounts and liquidity appear close to our baseline estimates (see Table 3.1, for example).

We move on to models (columns 2–4) that capture estimates for our term structure of credit risk-adjusted carbon premia with risk-neutral default probabilities. All models show that maturity, sector and carbon emissions are statistically significant, helping to explain corporate spreads. The correlation between carbon emissions and corporate spreads is positive, as in the core results. In terms of magnitude, the coefficients for bonds from non-energy-intensive companies are in the 1 to 3 range – close to our original results. For energy-intensive companies, the effect appears somewhat higher, with coefficients reaching a level of up to above 13 (10 in the original model). The twin term structures appear hump-shaped, nonetheless. In summary, our choice of default probability does not drive our results, which appear to hold in both the physical and risk-neutral worlds.

The second robustness check involves varying our liquidity measure. Chen et al (2007) use several measures of liquidity to show that the notion is priced in the cross section of corporate bond spreads. In our paper, we have chosen the absolute measure of Roll (see Appendix 2) as our preferred liquidity variable, given the simplicity of its computation and availability of the required data. In these alternative specifications, we explore whether using *observed* bid-ask spreads (as opposed to *synthetic* ones) affects our results. To this end, we gather close bid and ask yields for the corporate bonds in our sample from Bloomberg, and compute bid-ask spreads to re-run our key models. Table 6.2 summarises our estimates.

40 The theory behind the approach is that the equity of a firm can be viewed as a call option on the underlying value of the firm, with a strike price equal to the face value of the firm's debt. In brief: given a time series for the value of equity and liabilities for a particular firm, we can calibrate its corresponding asset values, the volatility of assets, and the probability of default. See Appendix 1 for details.

Preference channel models using risk-neutral PD

Table 6.1

	No emissions	Scope 1 emissions	Scope 1+2 emissions	Scope 1+2+3 emissions
Risk-neutral default probability (%)	3.09***	3.28***	3.31***	3.14***
	[0.39]	[0.38]	[0.38]	[0.37]
Duration	4.94***	2.45***	1.93***	1.48***
	[0.12]	[0.25]	[0.27]	[0.28]
Age	0.40***	0.36***	0.37***	0.37***
	[0.14]	[0.13]	[0.13]	[0.13]
Coupon	11.05***	9.28***	9.09***	8.91***
	[0.48]	[0.44]	[0.43]	[0.44]
ln(amount outstanding)	–2.83***	–2.91***	–2.92***	–2.94***
	[0.31]	[0.29]	[0.29]	[0.29]
Equity return volatility	31.08***	31.01***	30.99***	31.18***
	[0.92]	[0.91]	[0.91]	[0.90]
Liquidity	0.50***	0.47***	0.47***	0.47***
	[0.02]	[0.02]	[0.02]	[0.02]
Callable	–7.97***	–6.98***	–6.78***	–6.99***
	[1.00]	[0.95]	[0.94]	[0.94]
Non-energy-intensive x ln(emissions) x				
Maturity < 5 years		–0.57	2.66**	–10.21***
		[0.75]	[1.14]	[1.94]
Maturity 5-10 years		1.56**	4.76***	–8.10***
		[0.75]	[1.13]	[1.94]
Maturity 10-15 years		2.22***	5.48***	–7.31***
		[0.75]	[1.14]	[1.95]
Maturity 15-20 years		3.25***	6.53***	–6.25***
		[0.76]	[1.14]	[1.94]
Maturity 20-25 years		2.90***	6.29***	–6.46***
		[0.76]	[1.14]	[1.94]
Maturity 25-30 years		2.07***	5.61***	–7.00***
		[0.77]	[1.14]	[1.95]
Maturity > 30 years		2.24***	5.83***	–6.77***
		[0.79]	[1.16]	[1.96]
Energy-intensive x ln(emissions) x				
Maturity < 5 years		8.96***	12.60***	24.08***
		[2.58]	[3.18]	[3.95]
Maturity 5-10 years		11.00***	14.73***	25.84***
		[2.57]	[3.17]	[3.94]
Maturity 10-15 years		12.20***	16.03***	26.83***
		[2.56]	[3.17]	[3.93]
Maturity 15-20 years		13.33***	17.23***	27.88***
		[2.59]	[3.19]	[3.95]
Maturity 20-25 years		12.57***	16.55***	27.56***
		[2.55]	[3.16]	[3.94]
Maturity 25-30 years		11.52***	15.57***	26.81***
		[2.57]	[3.17]	[3.94]
Maturity > 30 years		12.85***	16.89***	28.02***
		[2.62]	[3.21]	[3.96]
Number of bonds	7,633	7596	7,596	7,596
Observations	266,133	263,771	263,886	263,886
R-squared	0.83	0.83	0.83	0.83

*** p<0.01, ** p<0.05, * p<0.1. Standard errors in brackets; clustered at the security level.

Sources: Bloomberg; Refinitiv; Trucost; authors' calculations.

Preference channel models using an alternative liquidity measure

Table 6.2

	No emissions	Scope 1 emissions	Scope 1+2 emissions	Scope 1+2+3 emissions
Default probability (%)	31.65***	31.91***	31.86***	31.59***
	[1.274]	[1.249]	[1.250]	[1.245]
Duration	5.236***	2.867***	2.375***	1.939***
	[0.116]	[0.249]	[0.267]	[0.277]
Age	0.569***	0.531***	0.535***	0.532***
	[0.134]	[0.126]	[0.125]	[0.126]
Coupon	10.51***	8.760***	8.585***	8.418***
	[0.470]	[0.420]	[0.417]	[0.420]
ln(amount outstanding)	-2.884***	-2.960***	-2.971***	-2.978***
	[0.313]	[0.288]	[0.286]	[0.289]
Equity return volatility	17.90***	17.87***	17.90***	18.05***
	[0.928]	[0.927]	[0.924]	[0.922]
Bid-ask spread	0.434***	0.407***	0.405***	0.405***
	[0.0169]	[0.0160]	[0.0159]	[0.0158]
Callable	-7.999***	-7.148***	-6.954***	-7.144***
	[0.987]	[0.933]	[0.926]	[0.927]
Non-energy-intensive x ln(emissions) x				
Maturity < 5 years		-1.088	2.449**	-5.760***
		[0.713]	[1.137]	[1.830]
Maturity 5-10 years		1.015	4.511***	-3.685**
		[0.706]	[1.129]	[1.825]
Maturity 10-15 years		1.637**	5.192***	-2.933
		[0.716]	[1.132]	[1.830]
Maturity 15-20 years		2.610***	6.187***	-1.926
		[0.719]	[1.135]	[1.826]
Maturity 20-25 years		2.263***	5.944***	-2.141
		[0.727]	[1.137]	[1.824]
Maturity 25-30 years		1.397*	5.219***	-2.727
		[0.729]	[1.139]	[1.826]
Maturity > 30 years		1.490**	5.368***	-2.567
		[0.758]	[1.153]	[1.835]
Energy-intensive x ln(emissions) x				
Maturity < 5 years		5.519**	7.900***	18.33***
		[2.349]	[2.901]	[3.591]
Maturity 5-10 years		7.557***	10.03***	20.07***
		[2.341]	[2.893]	[3.583]
Maturity 10-15 years		8.740***	11.30***	21.03***
		[2.337]	[2.891]	[3.580]
Maturity 15-20 years		9.818***	12.45***	22.02***
		[2.360]	[2.910]	[3.592]
Maturity 20-25 years		9.031***	11.73***	21.69***
		[2.321]	[2.878]	[3.582]
Maturity 25-30 years		8.007***	10.77***	20.94***
		[2.338]	[2.891]	[3.582]
Maturity > 30 years		9.279***	12.04***	22.11***
		[2.395]	[2.927]	[3.604]
Number of bonds	7,642	7599	7,599	7,599
Observations	266,241	263,682	263,797	263,797
R-squared	0.837	0.844	0.844	0.844

*** p<0.01, ** p<0.05, * p<0.1. Standard errors in brackets; clustered at the security level.

Sources: Bloomberg; Refinitiv; Trucost; authors' calculations.

This shows that: (1) the coefficient on observed bid-asks is statistically significant across all specifications, (2) it is positive, as expected; and (3) the order of magnitude is very close to that of our synthetic bid-ask measure. This is an encouraging outcome for the absolute measure of Roll as a liquidity proxy. Next, we focus on scope 1 and 2 emissions. With regard to their effect on spreads, the liquidity variable change induces minimal changes in the statistical power and magnitude of the tests. Our carbon premia results appear consistent with our choice of bid-ask spread.

6.2 Robustness checks for the risk channel

In testing the risk channel, we also consider the risk-neutral default probabilities that we compute on our own, to measure credit risk. The results are shown in Table 6.3.

The result showing that the risk channel is at work in both energy-intensive and non-energy-intensive sectors is robust to this alternative measure of default risk. The coefficients in front of scope 1 and scope 1+2 emissions for both energy and non-energy firms are positive and statistically significant. The coefficients in front of scope 1+2+3 emissions, however, lost significance. Comparing different sectors, coefficients in energy-intensive sectors are greater those in non-energy-intensive sectors, consistent with our main result in Section 4. Unsurprisingly, as risk-neutral default probabilities reflect both the probability of default *and* the default risk premium, the coefficients in front of both carbon emissions and control variables are larger.

Risk channel models using risk-neutral PD

Table 6.3

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Scope 1 emissions			Scope 1+2 emissions			Scope 1+2+3 emissions		
ln(emissions)	0.162	0.701***		0.210*	0.790***		-0.00716	0.141	
	[0.103]	[0.187]		[0.108]	[0.213]		[0.114]	[0.331]	
Non-energy-intensive x ln(emissions)			0.170*			0.215**			-0.0163
			[0.0973]			[0.101]			[0.110]
Energy-intensive x ln(emissions)			0.538***			0.585***			-0.0348
			[0.143]			[0.155]			[0.225]
ln(assets)	-0.882***	-1.31***	-0.91***	-0.94***	-1.41***	-0.97***	-0.73***	-0.652	-0.68***
	[0.124]	[0.278]	[0.113]	[0.135]	[0.301]	[0.122]	[0.136]	[0.424]	[0.130]
long-term debt/assets	9.072***	8.957***	8.902***	9.052***	9.070***	8.892***	9.241***	7.597***	8.926***
	[0.629]	[1.568]	[0.581]	[0.628]	[1.566]	[0.581]	[0.628]	[1.708]	[0.583]
Earnings/assets	-0.31**	-0.449	-0.37***	-0.303**	-0.421	-0.36***	-0.337**	-0.755	-0.35***
	[0.148]	[0.363]	[0.135]	[0.147]	[0.361]	[0.134]	[0.146]	[0.539]	[0.137]
Capital/assets	-2.341***	-3.629**	-2.39***	-2.46***	-3.859**	-2.51***	-2.24***	-4.041*	-2.29***
	[0.410]	[1.727]	[0.389]	[0.417]	[1.714]	[0.394]	[0.405]	[2.101]	[0.389]
Return on assets	-0.181***	-0.074**	-0.15***	-0.18***	-0.076**	-0.16***	-0.16***	-0.17***	-0.15***
	[0.0152]	[0.0296]	[0.0136]	[0.0153]	[0.0296]	[0.0136]	[0.0151]	[0.0420]	[0.0140]
Number of firms	2,295	374	2,670	2,295	374	2,670	2,324	345	2670
Observations	57,788	10,832	68,620	57,832	10,832	68,664	58,923	9,741	68664
R-squared	0.444	0.437	0.441	0.444	0.437	0.441	0.458	0.352	0.438

*** p<0.01, ** p<0.05, * p<0.1. Standard errors in brackets; clustered at the firm level.

Sources: Bloomberg; Refinitiv; Trucost; authors' calculations.

6.3 Robustness checks for total carbon premia

Our main analysis calculates total carbon premia in two steps. Indeed, premia through the preference and risk channels are computed, separately, in a first step, and then combined in a second step. However, as a robustness check, we can also estimate total carbon premia in one go. To do this, we simply need to swap our measure of default for a series of firm-level variables. This takes our regression model closer to those exploring the determinants of corporate spreads without isolating the risk channel. In other words, when doing this, our model looks less like Gilchrist and

Zakrajšek's (2012) and more like those found in Elton et al (2001), Campbell and Taksler (2003), and Chen et al (2007).

Concretely, we replace firm default probability by the natural logarithm of the firm's assets, its ratio of long-term debt to assets, its ratio of earnings to assets, its ratio of capital to assets and its return-on-assets. Table 6.4 shows our results. Overall, forecasting power does not suffer and the statistical significance on the triple interaction (sector, maturity, emissions) is preserved. Also, the coefficients on energy-intensive bonds appear slightly higher. Qualitatively, this alternative set of specifications does not alter our findings. Quantitatively, our estimated total impact is somewhat larger. For example, according to the estimates on Table 6.4, a halving in scope 1+2 carbon emissions would narrow spreads by around 8 and 18 basis points for non-energy-intensive firms and energy-intensive firms, respectively, at the belly of the curve. In contrast, our estimates in Section 5 suggest total effects of 5 and 10 basis points, respectively.

A simple model to compute total carbon premia (continues on next page) Table 6.4

	(1) No emissions	(2) Scope 1 emissions	(3) Scope 1+2 emissions	(4) Scope 1+2+3 emissions
ln(assets)	-26.44	-49.37**	-58.22***	-44.57**
	[20.44]	[19.73]	[19.58]	[20.25]
Long-term debt/assets	115.2***	117.9***	118.2***	116.5***
	[6.946]	[6.767]	[6.746]	[6.688]
Return on assets (%)	-1.739***	-1.737***	-1.761***	-1.700***
	[0.108]	[0.107]	[0.107]	[0.107]
Earnings/assets	-24.26***	-26.61***	-25.50***	-24.94***
	[3.514]	[3.286]	[3.237]	[3.265]
Capital/assets	12.31**	13.48***	13.14***	13.03***
	[4.974]	[4.914]	[4.919]	[4.917]
Duration	5.009***	2.404***	1.859***	1.381***
	[0.118]	[0.250]	[0.268]	[0.278]
Age	0.507***	0.475***	0.475***	0.474***
	[0.137]	[0.128]	[0.128]	[0.129]
Coupon	11.01***	9.187***	8.998***	8.827***
	[0.484]	[0.435]	[0.433]	[0.436]
ln(amount outstanding)	-2.959***	-3.044***	-3.055***	-3.062***
	[0.315]	[0.291]	[0.289]	[0.292]
Equity return volatility	29.09***	29.13***	29.12***	29.27***
	[0.874]	[0.857]	[0.853]	[0.850]
Liquidity	0.473***	0.445***	0.442***	0.442***
	[0.0178]	[0.0168]	[0.0167]	[0.0167]
Callable	-7.939***	-6.854***	-6.662***	-6.847***
	[1.008]	[0.955]	[0.947]	[0.949]

*** p<0.01, ** p<0.05, * p<0.1. Standard errors in brackets; clustered at the firm level. Sources: Bloomberg; Refinitiv; Trucost; authors' calculations. Table continues on next page.

A simple model to compute total carbon premia (continued)

Table 6.4

	(1) No emissions	(2) Scope 1 emissions	(3) Scope 1+2 emissions	(4) Scope 1+2+3 emissions
Non-energy-intensive x ln(emissions) x				
Maturity < 5 years		0.212	3.947***	-3.606*
		[0.849]	[1.252]	[1.977]
Maturity 5-10 years		2.366***	6.073***	-1.456
		[0.843]	[1.244]	[1.975]
Maturity 10-15 years		3.054***	6.829***	-0.628
		[0.851]	[1.246]	[1.980]
Maturity 15-20 years		4.090***	7.889***	0.449
		[0.854]	[1.248]	[1.978]
Maturity 20-25 years		3.806***	7.714***	0.281
		[0.861]	[1.250]	[1.979]
Maturity 25-30 years		3.000***	7.048***	-0.246
		[0.863]	[1.251]	[1.982]
Maturity > 30 years		3.131***	7.250***	0.00527
		[0.887]	[1.263]	[1.990]
Energy-intensive x ln(emissions) x				
Maturity < 5 years		9.830***	13.24***	21.55***
		[2.372]	[2.971]	[3.870]
Maturity 5-10 years		11.92***	15.43***	23.35***
		[2.364]	[2.964]	[3.862]
Maturity 10-15 years		13.11***	16.72***	24.34***
		[2.360]	[2.961]	[3.860]
Maturity 15-20 years		14.22***	17.91***	25.38***
		[2.378]	[2.975]	[3.871]
Maturity 20-25 years		13.48***	17.25***	25.16***
		[2.358]	[2.960]	[3.864]
Maturity 25-30 years		12.53***	16.37***	24.47***
		[2.364]	[2.964]	[3.862]
Maturity > 30 years		13.83***	17.66***	25.54***
		[2.418]	[2.996]	[3.884]
Number of bonds	7,396	7359	7,359	7,359
Observations	257,092	254,735	254,850	254,850
R-squared	0.83	0.838	0.838	0.838

*** p<0.01, ** p<0.05, * p<0.1. Standard errors in brackets; clustered at the firm level. Sources: Bloomberg; Refinitiv; Trucost; authors' calculations. The first part of this table is on the previous page.

7. Conclusions

In theory, corporate bond spreads represent compensation for bearing credit risk. In practice, they represent compensation for much more. As discussed, they encapsulate, among other things, compensation for the probability of default, lack of liquidity, higher volatility in firm value and, in the light of our results, increased firm-level pollution, as captured by greenhouse gas emissions.[41] The effect of firm-level

[41] Our result covers publicly traded companies exclusively.

emissions on corporate bond pricing has two aspects: firstly, regarding investor preferences; and secondly, regarding credit risk. We call these the preference and risk channels, respectively.

From a qualitative standpoint, the two channels arise for different reasons. First, investors may prefer holding debt issued by firms that are more environmentally friendly (vis-à-vis that of those that are not). This phenomenon evokes that of the liquidity premium, where on-the-run securities may be preferred to off-the-run ones, and compensation is due. Second, regardless of whether a firm is more favoured than another, some companies may be more exposed to risks during the transition to a low-carbon world. Carbon taxes, consumer preferences and technological change are only some of the factors that, if not planned for, could affect corporate financial health and therefore firm-level default risk.

From a quantitative standpoint, we find statistically significant evidence of both phenomena. In terms of economic significance, the impact is larger for energy-intensive firms in both channels. In addition, we find that the term structure of carbon premia – encapsulating both the preference channel and risk channels – is hump-shaped, with the largest premia at the belly of the curve (15–20 years). For a bond in this maturity bucket, which is issued by an energy-intensive firm, a halving of firm-level GHG emissions can reduce its spread by over 10 basis points.

Our results highlight the role of capital markets in the transition to net zero. Documenting the existence of a carbon premium provides evidence that investors differentiate between firms based on their carbon footprints. Such differentiation could incentivise firms to either reduce their GHG emissions or, less preferable from both the investor's and society's perspectives, to make it look as if they are doing so. Of course, any possibility of the latter underscores the need for strict disclosure standards or similar measures to protect investors and stakeholders from deception. In any case, whether the current size of the carbon premium can lead to a meaningful economic impact is a question yet to be investigated.

Our results also shed light on the financial stability implications of the decarbonisation transition. While we can take some comfort in the result that transition risks have been priced into corporate bonds, it remains to be seen if they have been priced in to a sufficient degree.

References

Agrawal, D, E Elton, M Gruber and D Mann (2001): “Explaining the rate spread on corporate bonds”, *Journal of Finance*, vol 56, pp 247–77.

Aswani, J, A Raghunandan, and S Rajgopal (forthcoming). "Are carbon emissions associated with stock returns? Reply", *Review of Finance*.

Amato, J and E Remolona (2003): “The credit spread puzzle”, *BIS Quarterly Review,* December 2003.

Amihud, Y and H Mendelson (1986): “Asset pricing and the bid-ask spread”, *Journal of Financial Economics*, vol 17, pp 223–49.

Bai, J, T Bali and Q Wen (2019): “Common risk factors in the cross-section of corporate bond returns”, *Journal of Financial Economics*, vol 131, pp 616–42.

Bai, J, R Golstein and F Yang (2020): “Is the credit spread puzzle a myth?”, *Journal of Financial Economics*, vol 137, pp 297–319.

Bharath, S, and T Shumway (2008): “Forecasting default with the Merton distance to default model”, *The Review of Financial Studies,* vol 21, no 3, pp 339–69.

Bolton, P, M Despres, L Pereira da Silva, F Samama and R Svartzman (2020): *The green swan: central banking and financial stability in the age of climate change,* Bank for International Settlements.

Bolton, P, and M Kacperczyk (2021a): “Do investors care about carbon risk?”, *Journal of Financial Economics*, vol 142, pp 517–49.

——— (2021b): “Global pricing of carbon-transition risk”, *NEBR Working Papers*, no 28510.

——— (forthcoming): "Are carbon emissions associated with stock returns? Comment", *Review of Finance*.

Bondioli, M, M Golberg, N Hu, C Li, O Maalaoui and H Stein (2021): *The Bloomberg Corporate Default Risk Model (DRSK) for Public Firms*, Bloomberg Finance LP.

Boss, M and M Scheicher (2002): “The determinants of credit spread changes in the euro area“, in *Market functioning and central bank policy*, *BIS Papers*, no 12, August, pp 181–99.

Briere, M and S Ramelli (2021): “Green sentiment, stock returns, and corporate behavior”, mimeo.

Busch, T, M Johnson and T Pioch (2018): “Consistency of corporate carbon emission data”, *University of Hamburg Report*, WWF Deutschland, Hamburg.

——— (2020): “Corporate carbon performance data: Quo vadis?”, *Journal of Industrial Ecology*, vol 26, no 1, pp 350–63.

Campbell, J and G Taksler (2003): “Equity volatility and corporate bond yields”, *Journal of Finance,* vol 58, pp 2321–49.

Carbone, S, M Giuzio, S Kapadia, J Krämer, K Nyholm and K Vozian (2021): "The low-carbon transition, climate commitments and firm credit risk", *ECB Working Papers*, no 2631.

Chava, S (2014): "Environmental externalities and cost of capital", *Management Science*, vol 60, pp 2223–47.

Chen, L, D Lesmond and J Wei (2007): "Corporate yield spreads and bond liquidity" *Journal of Finance,* vol 62, pp119–49.

Christopoulos, A (2020): "The Absolute Roll Measure".

Collin-Dufresne, P, R Goldstein and S Martin (2001): "The determinants of credit spread changes", *Journal of Finance*, vol 56, pp 2177–07.

Duffee, G (1998): "The relation between Treasury yields and corporate bond yield spreads", *Journal of Finance,* vol 53, pp 2225–41.

Duffie, D, L Saita and K Wang (2007): "Multi-period corporate default prediction with stochastic covariates", *Journal of Financial Economics*, vol 83, no 3, pp 635–65.

Duan, T, F Li and Q Wen (forthcoming). "Is carbon risk priced in the cross-section of corporate bond returns?", *Journal of Financial and Quantitative Analysi*s.

Ehlers, T, K Greiff and F Packer (2022): "The pricing of carbon risk in syndicated loans: which risks are priced and why?", *Journal of Banking & Finance*, vol 136, pp 106–80.

Elsenhuber, U and A Skenderasi (2019): "ESG investing: practical considerations for public investors", in *Evolving Practices in Public Investment Management: Proceedings of the Seventh Public Investors Conference*.

Felhütter, P and S Schaefer (2018): "The myth of the credit spread puzzle", *Review of Financial Studies,* vol 31, pp 351–67.

Ge, W and M Liu (2015): "Corporate social responsibility and the cost of corporate bonds", *Journal of Accounting and Public Policy*, vol 34, pp 597–624.

Giglio, S, B Kelly and J Stroebel (2021): "Climate finance", *Annual Review of Financial Economics*, vol 13, pp 15–36.

Gilchrist, S and E Zakrajšek (2012): "Credit spreads and business cycle fluctuations", *American Economic Review*, vol 102, no 4, pp 1692–720.

Goss, A and G Roberts (2011): "The impact of corporate social responsibility on the cost of bank loans", *Journal of Banking and Finance*, vol 35, pp 1794–810.

Greenhouse Gas Protocol (GGP)(2004): "Corporate Accounting and Reporting Standard, Revised Edition".

Gueant, O (2019): "Measuring liquidity on the corporate bond market", *Scientific Working Papers of the Autorité des Marchés Financiers* (AMF), March.

He, Z and K Milbradt (2014): "Endogenous liquidity and defaultable bonds", *Econometrica*, vol 82, no 4, pp 1443–508.

Heck, S (2021): "Corporate bond yields and returns: a survey", *Financial Markets and Portfolio Management*, Springer; Swiss Society for Financial Market Research, vol 36, no 2, pp 179–201.

Huynh, TD and Y Xia (2021): "Climate change news risk and corporate bond returns", *Journal of Financial and Quantitative Analysis*, vol 56, no 6, pp 1985-2009.

Intergovernmental Panel on Climate Change (IPCC)(2022): "Summary for policymakers", in P Shukla, J Skea, R Slade, A Al Khourdajie, R van Diemen, D McCollum, M Pathak, S Some, P Vyas, R Fradera, M Belkacemi, A Hasija, G Lisboa, S Luz and J Malley (eds), Climate change 2022: mitigation of climate change, Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change, Cambridge University Press, Cambridge, United Kingdom and New York.

Merton, R (1974): "On the pricing of corporate debt: the risk structure of interest rates", *Journal of Finance*, vol 29, no 2, pp 449–70.

Painter, M (2020): "An inconvenient cost: The effects of climate change on municipal bonds", *Journal of Financial Economics*, vol 135, pp 468–82.

Polbennikov, S, A Desclee, L Dynkin and A Maitra (2016): "ESG ratings and performance of corporate bonds", *The Journal of Fixed Income,* vol 26, no 1, pp 21–41.

Roll, R (1984): "A simple implicit measure of the effective bid-ask spread in an efficient market", *Journal of Finance*, vol 39, no 4, pp 1127–39.

Rossi, M (2014): "Realized volatility, liquidity and corporate yield spreads", *Quarterly Journal of Finance,* vol 4, no 1.

Seltzer, L, L Starks and Q Zhu (2022): "Climate regulatory risks and corporate bonds", Federal Reserve Bank of New York, *Staff Reports*, no 1014, April.

Appendix 1. Estimating risk-neutral default probabilities

It is possible to estimate default probabilities using Merton's structural model (1974). We start from the assumption that the total value of the firm V (its assets) follows a geometric Brownian motion:

$$dV = \mu_V V dt + \sigma_V V dW$$

where μ_v is the expected return on the value of the firm, σ_V is the volatility of the firm's value and dW is an increment of the standard Weiner process. Next, we must make an assumption about the firm's capital structure. It is assumed that the firm has issued D amount of a single zero-coupon bond of T years maturity.

These assumptions imply that the value of the firm's equity, which we denote E, can be viewed as a call option on the underlying value of the firm V, with a strike price equal to the face value of the firm's debt D and a time to maturity of T. According to the Black-Scholes pricing formula, the value of the firm's equity (the "put option") is given by:

$$E = V\Phi(d_1) - e^{-rT} D\Phi(d_2)$$

where $\Phi(\cdot)$ is the cumulative standard normal distribution function and r the risk-free rate, which is used to continuously discount the value of the debt. Furthermore:

$$d_1 = \frac{ln(V/D) + (r + 0.5\sigma_V^2)T}{\sigma_V^2\sqrt{T}}, d_2 = \delta_1 - \sigma_V\sqrt{T}$$

This way, the value of the firm's equity depends on the total value of the firm and time, which allows us to relate the volatility of the firm's value σ_V to the volatility of its equity σ_E. From Ito's Lemma, and given that under this option pricing framework $\delta E/\delta V = \Phi(d_1)$, we can derive this relationship as:

$$\sigma_E = \left(\frac{V}{E}\right)\Phi(d_1)\sigma_V$$

The inputs to the Merton model are therefore the value of equity, the value of debt and the volatility of equity. Naturally, because a company's debt structure is more complex than the aforementioned zero-coupon bond, we assume that the debt threshold is somewhere between the face value of the short-term debt (D_{ST}) and long-term debt (D_{LT}). Concretely:

$$D = D_{ST} + 0.5D_{LT}$$

In addition, we use a horizon T of five years in total, which matches the longest horizon available for Bloomberg's default probabilities. Firm data for the model is collected from S&P Capital IQ as discussed in section 2.2. For the implementation and, as in Gilchrist and Zakrajšek (2012), we use an iterative procedure proposed by Bharath and Sumway (2008), which addresses large swings in estimated volatility for the firm's value σ_E. For a time series of inputs, the outputs of the model are a time series of σ_V and V. We can then use these to compute firm-specific default probability PD as:

$$PD = \Phi\left(-\frac{ln(V/D) + (\mu_V - 0.5\sigma_V^2)}{\sigma_V}\right)$$

Appendix 2. The absolute Roll measure of illiquidity

The general theory behind the liquidity premium originated with Amihud and Mendelson (1986) and, since at least Chen et al (2007), liquidity has been documented as an important determinant of credit spreads. In their original work, realised bid–ask spreads appear as one of the central measures for gauging bond market illiquidity. However, the absence of *observable* (as opposed to *quoted*) bid-ask spreads for corporate bonds is an issue for the data gathering process (see Gueant (2019) for example).

In its stead, interesting alternatives have been proposed. A theoretically attractive one is the Roll (1984) measure, which allows us to compute a theoretical or *effective* bid-ask, based solely on daily closing price data. In brief, if p_t is the end-of-day price for a bond, we can compute its effective bid-ask spread λ as:

$$\lambda = 2\sqrt{-Cov(\Delta P_t, \Delta P_{t+1})}$$

where $Cov(\Delta P_t, \Delta P_{t+1})$ is the autocovariance of price changes. From this expression, it is easy to see that a complex number is derived when $Cov(\Delta P_t, \Delta P_{t+1}) > 0$. Alternative versions of the measure have been proposed to address the issue – for example, by dropping those observations where the bid-ask spread could potentially be negative. However, these methods lead to gaps in the data. An approach which seeks to preserve the amount of input data available is proposed by Christopoulos (2020) and dubbed the *absolute* Roll measure.

The absolute Roll measure $\hat{\lambda}$ is given by:

$$\hat{\lambda} = 2\sqrt{|-Cov(\Delta P_t, \Delta P_{t+1})|}$$

which leads to a strictly non-negative bid-ask spread applicable to all traded securities that are limited to closing price information. Given the availability of closing price data for our bond sample, we favour the use of this measure in our model of corporate spreads.

To procure monthly data (as our panel regression requires), we follow these steps:

1. For each bond j, we gather all daily closing price data available for month t. Assuming 20 trading days per month, this is a time series of daily prices $\{p^j_{t/20}, p^j_{2t/20}, \ldots, p^j_t\}$.

2. We compute the absolute Roll measure as the autocovariance of this process.

3. We store this computation as the effective bid-ask for month t, $\hat{\lambda}_t$ for each bond j.

4. We repeat this process for the following month, across all bonds.

Appendix 3. Regression analysis with carbon intensities

Our main results are based on analysing a firm's carbon footprint using carbon emissions, which is consistent with Bolton and Kacperczyk (2021a), who link the carbon premium in stock returns to carbon emissions. This choice is also aligned with regulatory frameworks, such as climate stress tests, which tend to focus on activities with high levels of emissions.

To further test the robustness of our findings, we have repeated our core analysis using another common measure of a firm's carbon footprint: carbon emission intensity, which is defined as the ratio of carbon emissions to revenue. This metric has been used in prior studies, such as Ehlers et al (2022) and Duan et al (forthcoming), who investigate the relationships between syndicated loan spreads and corporate bond returns, respectively, and carbon emission intensities.[42]

Tables A3.1 and A3.2 show the results of our analyses using carbon emission intensities in place of carbon emissions. They serve as analogues for Tables 3.4 and 4.2, respectively, in the body of the paper. Our findings suggest that carbon emission intensities are priced into corporate bond spreads through both preference and risk channels. However, we observe that the maturity component (and thereby, the term structure of carbon premia) is statistically significant only in non-energy-intensive sectors.

The preference channel: results with carbon intensities — Table A3.1

	(1) Scope 1 emissions	(2) Scope 1+2 emissions	(3) Scope 1+2+3 emissions
Non-energy-intensive x ln(intensity) x			
Maturity < 5 years	0.143	3.460***	1.209
	[0.820]	[1.236]	[2.306]
Maturity 5-10 years	4.412***	9.274***	7.096***
	[0.808]	[1.208]	[2.286]
Maturity 10-15 years	5.235***	10.57***	9.226***
	[0.913]	[1.252]	[2.308]
Maturity 15-20 years	6.309***	13.11***	12.27***
	[0.922]	[1.279]	[2.311]
Maturity 20-25 years	4.602***	11.36***	11.34***
	[0.911]	[1.301]	[2.318]
Maturity 25-30 years	0.99	7.599***	9.271***
	[0.871]	[1.281]	[2.317]
Maturity > 30 years	0.642	7.430***	9.608***
	[1.031]	[1.366]	[2.369]
Energy-intensive x ln(intensity) x			
Maturity < 5 years	-3.457*	-5.065**	-3.848
	[2.024]	[2.389]	[4.041]
Maturity 5-10 years	0.316	-0.871	0.274

[42] See Bolton and Kacperczyk (forthcoming) and Aswani et al (forthcoming) for a debate regarding which metric better measures a firm's carbon footprint: carbon emissions or carbon intensities.

	[2.004]	[2.367]	[4.019]
Maturity 10-15 years	1.953	1.225	2.612
	[2.025]	[2.390]	[4.020]
Maturity 15-20 years	3.782*	3.436	5.083
	[2.097]	[2.441]	[4.045]
Maturity 20-25 years	1.315	1.255	3.878
	[1.992]	[2.368]	[4.019]
Maturity 25-30 years	-1.803	-1.545	1.749
	[2.010]	[2.382]	[4.032]
Maturity > 30 years	1.18	1.432	4.578
	[2.397]	[2.671]	[4.132]
Number of bonds	7599	7599	7599
Observations	263,682	263,682	263,682
R-squared	0.84	0.843	0.844

*** p<0.01, ** p<0.05, * p<0.1. Standard errors in brackets; clustered at the security level. Coefficients on other variables omitted for brevity. Sources: Bloomberg; Refinitiv; Trucost; authors' calculations.

The risk channel: results with carbon intensities

Table A3.2

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Scope 1 emissions			Scope 1+2 emissions			Scope 1+2+3 emissions		
ln(inten.)	0.07***	0.11***		0.09***	0.12***		0.00	0.14***	
	[0.01]	[0.02]		[0.02]	[0.03]		[0.04]	[0.03]	
Non-energy-intensive x ln(inten.)			0.07***			0.09***			0.14***
			[0.01]			[0.02]			[0.03]
Energy-intensive x ln(inten.)			0.11***			0.12***			0.00
			[0.02]			[0.03]			[0.04]
ln(assets)	-0.14***	-0.15***	-0.14***	-0.14***	-0.15***	-0.14***	-0.14***	-0.14***	-0.14***
	[0.01]	[0.02]	[0.01]	[0.01]	[0.02]	[0.01]	[0.03]	[0.01]	[0.01]
long-term debt/assets	1.56***	1.62***	1.55***	1.54***	1.63***	1.54***	1.84***	1.54***	1.57***
	[0.09]	[0.22]	[0.08]	[0.08]	[0.22]	[0.08]	[0.26]	[0.08]	[0.08]
Earnings/assets	-0.12***	-0.19***	-0.14***	-0.12***	-0.18**	-0.13***	-0.13	-0.13***	-0.13***
	[0.02]	[0.07]	[0.02]	[0.02]	[0.07]	[0.02]	[0.09]	[0.02]	[0.02]
Capital/assets	-0.45***	-0.31	-0.44***	-0.50***	-0.33	-0.49***	-0.38	-0.50***	-0.50***
	[0.05]	[0.28]	[0.04]	[0.05]	[0.28]	[0.05]	[0.25]	[0.05]	[0.04]
Return on assets	-0.05***	-0.03***	-0.04***	-0.05***	-0.03***	-0.04***	-0.04***	-0.05***	-0.04***
	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
Number of firms	2,443	388	2,831	2,443	388	2,831	365	2466	2831
Observations	121,112	19,650	140,762	121,208	19,650	140,858	18,018	122,840	140858
R-squared	0.49	0.55	0.5	0.49	0.55	0.5	0.57	0.49	0.5

*** p<0.01, ** p<0.05, * p<0.1. Standard errors in brackets; clustered at the firm level. Sources: Bloomberg; Refinitiv; Trucost; authors' calculations.