Machine learning and micro-prudential climate stress testing

Christian Haas*[1], Karol Kempa[1], Ulf Moslener[1], and Sebastian Rink[1]

[1]Frankfurt School of Finance & Management

This version: November 2025

Abstract

Rising climate-related transition risks demand more robust risk management tools. This paper introduces a novel climate stress testing framework that applies machine learning methods to examine firm-level carbon price shocks. The framework extends existing approaches by explicitly modelling small and medium-sized enterprises (SMEs) and incorporating prediction uncertainty via conformal prediction. We estimate that a EUR 100 carbon price shock nearly doubles the number of loss-making firms and triggers an increase in loan defaults by 7.9%. These effects are significantly amplified by shock size. SMEs are disproportionately affected, whereas large corporations show greater resilience. Substantial uncertainty remains due to limited emissions reporting, particularly among SMEs. Overall, the results underscore both the value of machine learning for enhancing the granularity and reliability of climate stress tests and the need for more comprehensive firm-level climate disclosure.

JEL-Codes: C53, C55, G33, Q54.

Keywords: Stress test, Machine learning, Climate risk, Shock, Prediction uncertainty, Climate data.

The authors acknowledge financial support from the research fund Klimaschutz und Finanzwirtschaft (KlimFi) of the German Federal Ministry of Education and Research (Grant No. 01LA2210A).

* Corresponding author: c.haas@fs.de.

I. Introduction

The links between climate change and the financial system have gained increasing attention. On the one hand, the financial system may play a role in achieving climate targets (Battiston et al., 2021). On the other hand, the financial system is increasingly exposed to the risks arising from climate change. These risks include both risks due to physical impacts of climate change and risks stemming from the transition towards a net-zero economy, where the latter are of particular concern (Acharya et al., 2023; Monasterolo, 2020). Limiting global warming, as set out in the Paris Agreement, will require substantially more ambitious climate policies (Nordhaus, 2018). Although carbon pricing has gained importance worldwide, it still only covers 24% of global greenhouse gas (GHG) emissions as of 2024.[1] At the same time, the social costs of carbon have increased over time (Tol, 2023). Abrupt increases in the ambition of mitigation regulation could cause economic costs for the real economy that could be transmitted to the financial sector via defaults and drops in asset prices (Bolton and Kacperczyk, 2023; Seltzer et al., 2022; Kempa and Moslener, 2022; Capasso et al., 2020; de Bandt et al., 2025), potentially threatening financial stability (Campiglio et al., 2023). One approach that has been applied to investigate this threat is climate stress testing. A key obstacle to micro-prudential stress test is the limited availability of climate-related firm-level data, particularly on GHG emissions. This paper addresses this issue by introducing machine learning methods to improve climate stress testing with an application to the micro-prudential level in Europe.

In general, stress tests are a risk management instrument in the financial sector and have been applied by regulatory authorities to test the stability of the financial system in the event of a major economic disruption. Stress tests can be divided into microprudential stress tests, typically performed by the macroprudential authority or financial institutes themselves, and macroprudential stress tests, focusing on the resilience of the financial

[1] Source: World Bank Climate Pricing dashboard, available at `https://carbonpricingdashboard.worldbank.org/compliance/coverage` (last accessed on 7 April 2025).

system as a whole. More recently, such stress tests have been conducted to investigate climate-related risks by both central banks, e.g., the European Central Bank (ECB, 2022) or the Bank of England (Bank of England, 2022), but also increasingly by researchers (Nguyen et al., 2023; Battiston et al., 2017; Nieto, 2019). Furthermore, previous studies have proposed new methodological approaches, such as forward-looking risk measures, the creation and use of scenarios, or financial networks, that could be applied within climate stress tests (Baer et al., 2022, 2023; Battiston et al., 2021; Battiston and Martinez-Jaramillo, 2018; Battiston et al., 2019; Emambakhsh et al., 2023; Reinders et al., 2023; Jung et al., 2021; Kainth et al., 2024; Koberle et al., 2021; Xu et al., 2024). Reinders et al. (2025) provide a recent overview of climate stress testing exercises.

This paper advances the field (Acharya et al., 2023; Battiston et al., 2017; Baer et al., 2022) by addressing key limitations identified in recent reviews (Acharya et al., 2023; Reinders et al., 2025; Koberle et al., 2021), in particular by predicting firm-level GHG emissions and hence increasing firm coverage. We develop a micro-prudential stress testing framework that integrates machine learning techniques. Our stress testing framework comprises three steps. First, we train a machine learning model using the framework proposed by Haas et al. (2025) on a large dataset of firms, including small and medium-sized enterprises (SMEs). Then we estimate firm-level Scope 1 emissions using the retrained model on financial and operational data from the BvD Orbis database. This approach avoids reliance on incomplete vendor data and enables comprehensive coverage of indebted firms. Second, using 2022 data, we analyse the effect of a EUR 100 per tonne carbon price shock, consistent with previous regulatory exercises (ECB, 2022), as the base scenario. We analyse the shock using 2022 data, considering all indebted firms in climate policy-relevant sectors (Battiston et al., 2017).[2] The shock increases firms' costs proportional to their emissions, and we assume no pass-through to consumers. Based on accounting definitions, the profit and loss statements and balance sheets are adjusted accordingly. In addition to this base scenario, we simulate a range of carbon price shocks (EUR 1–500 per

[2]As we are interested in the effect on credit loss, we do not consider firms without debt capital.

tonne of CO_2)[3]. Third, to assess the impact of shocks on defaults, we develop a binary machine learning classification model. The model incorporates firm-level financial data and macroeconomic variables, with lagged features that capture temporal dynamics.

We find that a carbon price shock of EUR 100 on top of the existing level stresses the financial viability of firms, resulting in a loss of asset value and, in particular, adversely affecting profit margins and doubling the number of firms incurring losses. In addition, such a shock induces the default of almost 8,000 firms, resulting in a substantial increase in defaulting loan volume by 7.9%. All these adverse impacts increase with the size of the carbon price shock. Adverse impacts are particularly strong for smaller firms, whereas very large corporations tend to be more resilient to such shocks. The sectoral analysis reveals substantial heterogeneity across sectors, with some sectors, such as mining or transport, experiencing particularly high adverse impacts.

We further explore the role of firm heterogeneity by comparing our approach based on firm-level emissions with an alternative approach using averages of industry-level emissions. Using the latter and hence ignoring firm heterogeneity within a sector, tends to underestimate the adverse effects of a carbon price shock. The magnitude of this potential underestimation is highly sector specific. Industry averages are a reasonable imputation method for low-emission sectors. Firm-specific emissions are more important in high-emission sectors due to high firm heterogeneity, such as construction or manufacturing. Finally, we show that default prediction results are sensitive to uncertainty in the emission estimation, in particular because of the low availability of data, which may be further aggravated by ongoing policy role backs of ESG disclosures.

Our contributions are threefold. First, we use machine learning to predict firm emissions for further use within climate stress tests. Based on the predicted direct (Scope 1) emission data, we can compute the additional costs associated with the carbon price shock

[3] Tol (2023) argues that the social cost of carbon has gone up to more than 500 US$ per tonne of CO_2, depending on the discount rate.

by multiplying the carbon price by firm-level CO_2 levels. We do not have to rely on approximation methods for the effect of a carbon price shock such as value-added taxes (VAT) in Alogoskoufis et al. (2021). Hence, the stress test is more precise, as the impact of a carbon price shock on the firm's likelihood of default depends on their CO_2 emissions. By introducing conformal emission prediction intervals, we can assess the effect of uncertainty in the emission prediction on the stress-testing results. This contribution is particularly relevant, as recent regulatory rollbacks, such as Europe's Omnibus process (European Commission, 2025) and the SEC's decision to halt new sustainability disclosure rules (Pinedo et al., 2025), suggest firm-level sustainability data will remain limited in major economies.

Second, we can extend the scope of climate stress tests by using a large dataset of European firms. Previous climate stress tests focus on major banks in the EU (Battiston et al., 2017; ECB, 2022), banks in the US (Nguyen et al., 2023), or the banking sector in other individual countries (Jung et al., 2021; Reinders et al., 2023; Grippa and Mann, 2021; Vermeulen et al., 2019). In contrast to these studies, we do not focus on banks, but on the whole portfolio of European firms with available financial information. Most importantly, we also include SMEs, which represent a substantial share of banks' credit exposures. Our dataset comprises almost 1.4 million unique EU firms. Moreover, stress scenarios typically apply static shock levels, limiting their ability to capture the full range of potential impacts (ECB, 2022; Bank of England, 2022), while we analyse carbon price shocks of EUR 1–500 per tonne of CO_2. By providing a range of stress levels and a comprehensive coverage of firm types, we contribute to ongoing efforts to integrate climate-related risks into financial supervision and risk management.

Third, we also apply machine learning to predict firm defaults. In contrast to previous papers with model-based approaches, e.g. using the Merton model to estimate the probability of default (Nguyen et al., 2023; Reinders et al., 2023), we use a model-free approach. Instead of modelling a specific relationship between firm default and a theory-based selection of explanatory variables, the machine learning approach combines financial economic theory with information from the data to approximate the functional form. This method optimises the bias-variance trade-off, which results in a better prediction accuracy.

The paper is structured as follows. Section II presents the methodology and data. Section III reports and discusses the main results and Sections IV & V present additional analyses. Section VI concludes.

II. Methodology

A. General Framework

We employ a bottom-up, firm-level framework to assess the exposure of European firms to transitory climate risk. The analysis considers counterfactual scenarios with exogenous increases in the price of Scope 1 firm-level GHG emissions. We proceed in three steps. First, to achieve broad coverage of firms and to address in particular the limited availability of reported emission data for medium and non-listed firms[4], we build a supervised machine-learning framework to estimate firm-level Scope 1 GHG emissions. Second, we consider carbon-price increases ranging from 1 to 500 EUR per tCO_2e and translate these increases – hereafter, 'shocks' – into accounting-consistent firm financials. This is done by mapping the shock into firms' cost structures and propagating it through the profit and loss account, the balance sheet, and cash-flow statements to obtain scenario-adjusted financials. Third, we build a supervised machine-learning classifier framework that uses the resulting shock-adjusted financials, together with firm fundamentals, to predict binary firm-default events. Comparing no-shock and shock scenarios yields an estimate of the (marginal) effect of carbon-price shocks on firm default (risk) and banks' portfolio risk. Figure 1 summarises the workflow of the conceptual framework. Details on the

[4]Firm size is defined as: *Very large (XL)*: at least one of revenue $\geq$ 100 m EUR, assets $\geq$ 200 m EUR, employees $\geq 1,000$, or listed; *Large (L)*: at least one of revenue $\geq$ 10 m EUR, assets $\geq$ 20 m EUR, or employees $\geq$ 150, not XL; *Medium (M)*: at least one of revenue $\geq$ 1 m EUR, assets $\geq$ 2 m EUR, or employees $\geq$ 15, not L or XL.

steps, including the model frameworks and estimation procedures are provided in subsections B-D. All workflows are implemented in Python 3.11.11 using scikit-learn (Pedregosa et al., 2011) for preprocessing steps and construction of model pipelines. Finally, Section E presents the data and variables.

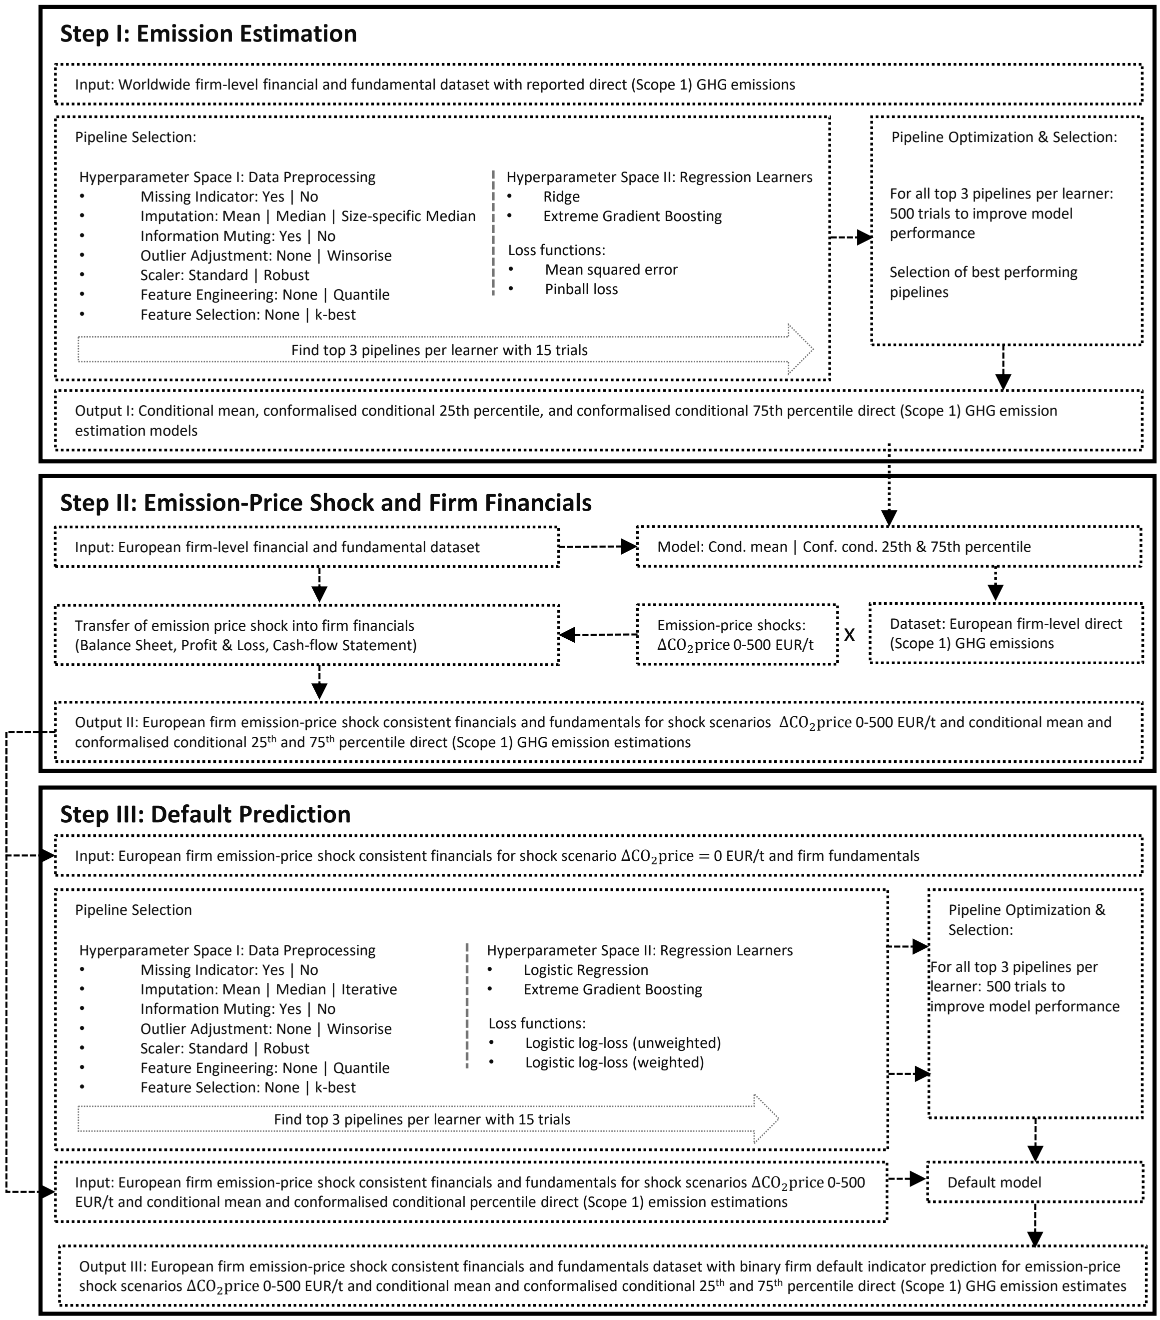

Figure 1. Climate Stress Test Framework

B. Emission Estimation Model

We estimate the conditional mean and the 25th and 75th conditional quantiles of firms' annual emissions. The conditional mean serves as the central emission estimate for subsequent analyses, while the quantiles—here conformalised versions of classical quantile estimates—provide measures of estimation uncertainty. The emission–estimation framework utilises a global annual firm–level dataset with a broad set of features (covariates), including fundamentals (i.a., sector, country, and firm–size indicators) as well as financial variables and ratios (see Section E). The target (dependent) variable is annual Scope 1 CO_2-equivalent (CO_2e) emissions in tonnes.[5] We only use reported emissions in training and evaluation.

For the conditional-mean estimation, we apply the framework by Haas et al. (2025) and construct heterogeneous machine-learning pipelines. The framework enables the estimation of firm-level sustainability data in a domain-agnostic and replicable manner. Each pipeline follows the same sequence of steps: (i) missing-data indicator, (ii) imputation, (iii) information muting, (iv) outlier adjustment, (v) scaling, (vi) feature engineering, (vii) feature selection, and (viii) regression. At each step, multiple options are available. We form a large set of candidates by selecting exactly one option per step, enabling a systematic search for the configuration best suited to the task. For missing-data indicators, the options include adding missing-value flags for numeric features or omitting them. Imputation uses either full-sample mean/median replacement or firm-size–specific medians. Optional information muting removes firm-size dummies during training to avoid mechanical correlations in sparse labels. Outlier handling is either not applied or uses winsorisation at the 5th/95th percentile levels. Scaling uses a robust scaler to 'normalize' the distribution. To cope with potential covariate shift, one scaler option is pre-fitted on the full dataset for consistent application at inference. Additional optional steps include simple feature engineering adding squared terms of numeric features and feature selection

[5] For both emission model training and descriptive displays, we apply the transformation $\log(\text{Scope 1 } CO_2e_t + 1)$ to accommodate zeros and reduce skewness.

via a fixed k-best filter, where we set $k = 25$. The final regression stage uses linear and non-linear learners: ridge regression (Hoerl and Kennard, 1970) for stable extrapolation, and gradient-boosted trees (XGBoost, CatBoost) for capturing flexible nonlinearities and interactions (Chen and Guestrin, 2016; Prokhorenkova et al., 2018).

To fit and select pipelines with statistically valid evaluation, we partition the data by firm identifier so that each firm (and all its annual observations) appears in exactly one set. Firms are randomly assigned to training (65%), calibration (25%), and test (10%) sets. The training set is used to fit and select conditional-mean pipelines and quantile pipelines. The calibration set is employed for conformal adjustment of quantiles. Evaluation of fitted conditional-mean pipelines is based on the joint set of calibration set and test set, while interval coverage based on estimations from quantile pipelines is assessed on the test set only. Model selection proceeds in three stages. First, all candidate pipelines are fitted and evaluated in 15 iterations of Bayesian hyperparameter optimisation using Optuna (Akiba et al., 2019), with grouped ten-fold cross-validation (grouped by firm) and minimisation of mean out-of-fold MSE. This screening allows to rank the different pipeline configurations by predictive performance. Second, the three best-performing pipelines per regressor are selected and undergo 500 additional optimisation iterations. In each iteration, hyperparameters are proposed as in the first step, model parameters are fitted, and performance is evaluated as before. Early stopping terminates trials that underperform running medians of mean out-of-fold MSE across pipelines per regressor. In the next stage, the final hyperparameters are set to those of the best-performing configuration from the second stage, and model parameters are refitted on the full training set. The resulting model maps firm features to a point estimate of annual emissions.

On the hold-out test set, non-linear boosted-tree pipelines outperform linear ridge pipelines in terms of MSE (see Table I, panel a). RMSEs for the best pipelines are 1.7619 for boosted trees versus 3.5376 for ridge. For firms with reported emissions, the conditional-mean model captures the emissions distribution well across all years (Figure 2, left panel)

Table I. Performance of Point-/Interval-GHG Emission Estimation Models

Target	(a) Conditional mean		(b) Conformal quantile interval		
Rank	RMSE	MAE	Emp. coverage	Avg. width	Var(width)
1st-best linear[a]	3.5376	1.7596	51.88%	3.732492	15.844285
1st-best non-linear[b]	1.7619	1.2363	50.87%	3.003459	0.719601
2nd-best linear[c]	3.5377	1.7592	51.90%	3.325559	15.986328
2nd-best non-linear[d]	1.7627	1.2375	50.88%	3.004581	0.721762

[a]Conditional mean (MSE loss) pipeline: no missing-indicator; size-specific median imputation; muting of firm-size dummies; no outlier removal; pre-fitted robust scaler (inference sample); no feature engineering, no additional transformations, no feature selection; ridge regressor. Conformal interval (pinball loss + split-conformal) pipeline: no missing-indicator; size-specific median imputation; no muting; no outlier removal; pre-fitted robust scaler (inference sample); no feature engineering, no additional transformations, no feature selection; ridge regressor.

[b]Conditional mean (MSE loss) pipeline: missing-indicator augmentation (numeric); size-specific median imputation; muting of firm-size dummies; no outlier removal; robust scaler; no feature engineering, no additional transformations, no feature selection; XGBoost regressor. Conformal interval (pinball loss + split-conformal): no missing-indicator; size-specific median imputation; muting of firm-size dummies; no outlier removal; robust scaler; no feature engineering, no additional transformations, no feature selection; XGBoost regressor.

[c]Conditional mean (MSE loss) pipeline: no missing-indicator; size-specific median imputation; no muting; no outlier removal; robust scaler; no feature engineering, no additional transformations, no feature selection; ridge regressor. Conformal interval (pinball loss + split-conformal) pipeline: no missing-indicator; size-specific median imputation; muting of firm-size dummies; no outlier removal; robust scaler; no feature engineering, no additional transformations, no feature selection; ridge regressor.

[d]Conditional mean (MSE loss) pipeline: missing-indicator augmentation (numeric); size-specific median imputation; muting of firm-size dummies; no outlier removal; pre-fitted robust scaler (inference sample); no feature engineering, no additional transformations, no feature selection; XGBoost regressor. Conformal interval (pinball loss + split-conformal) pipeline: missing-indicator augmentation (numeric); size-specific median imputation; no muting; no outlier removal; pre-fitted robust scaler (inference sample); no feature engineering, no additional transformations, no feature selection; XGBoost regressor.

and specifically for 2022, the stress-test year (Figure 2, right panel). The 2022 distribution appears slightly left-shifted across firm sizes when considering the full sample of reporting and non-reporting firms (Figure 3). In addition to firm-level conditional-mean estimates, we aggregate estimated emissions to NACE-2 divisions (2-digit level) and equally reallocate them to constituent firms. These sectoral average emission estimates allow us to assess the relevance of firm-specific emission estimations in contrast to the use of industry averages to predict firm defaults.

Uncertainty quantification builds on the conditional-quantile pipelines, which are fitted, evaluated, and selected using the same pipeline architecture and training steps as described above for the conditional-mean estimation pipelines, but trained using the pin-

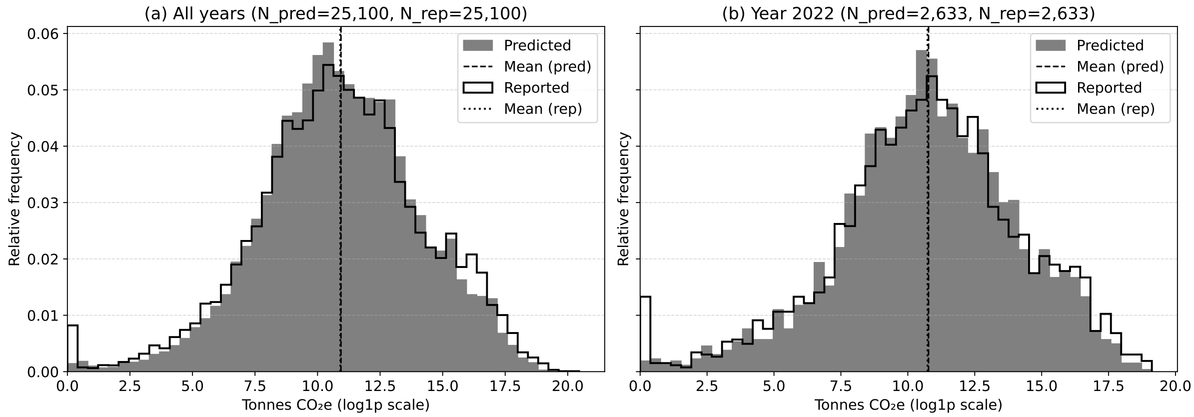

Notes: This figure presents the distribution of conditional mean predictions of firm-level Scope 1 GHG emissions and the distribution of actual reported emissions for a sub-sample of firms for which reported data are available. For the prediction, a XGBoost model is used. Panel (a) presents the distribution for all years in the sample and panel (b) for the year 2022, which is used in the stress testing exercise.

Figure 2. Distribution of Predicted vs. Reported Scope 1 GHG Emissions for Reporting Firms

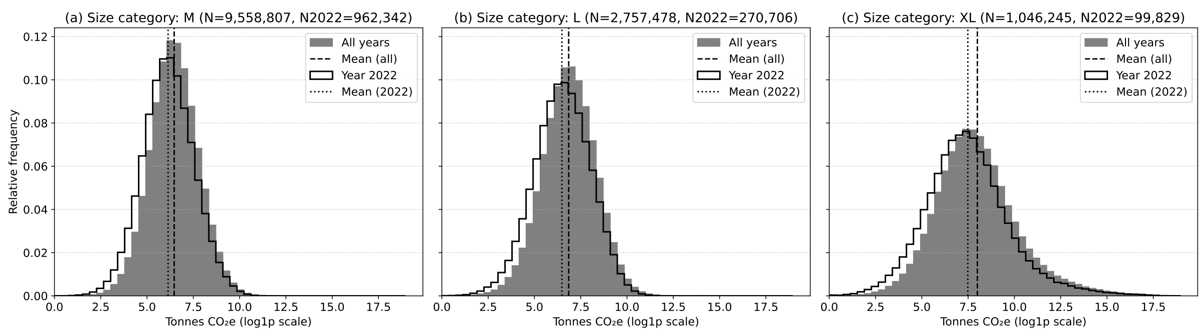

Notes: This figure shows the distribution of estimated Scope 1 GHG emissions for different firm sizes: panel (a) for M-sized firms, panel (b) for L-sized firms, and panel (c) for XL-sized firms. Firm size is defined as follows: *Very large (XL)*: at least one of revenue ≥ 100 m EUR, assets ≥ 200 m EUR, employees $\geq 1,000$, or listed; *Large (L)*: at least one of revenue ≥ 10 m EUR, assets ≥ 20 m EUR, or employees ≥ 150, not XL; *Medium (M)*: at least one of revenue ≥ 1 m EUR, assets ≥ 2 m EUR, or employees ≥ 15, not L or XL. Each panel shows the distribution for all years as well as for the year 2022, which is used in the stress testing exercise. The underlying prediction model is XGBoost.

Figure 3. Distribution of Predicted Scope 1 GHG Emissions for Reporting Firms across Firm Sizes

ball loss function rather than mean squared error in the regressor algorithm. The fitted pipeline-configurations with minimal mean out-of-fold pinball loss are used to obtain raw point estimates for the 25th and 75th conditional quantiles. Based on Romano et al. (2019), the raw quantile estimates are then adjusted via split-conformalisation on the calibration set. Tail-specific nonconformity scores are used to adjust the raw quantiles to yield prediction intervals $[\tilde{q}_{0.25},\, \tilde{q}_{0.75}]$, where $\tilde{q}_\alpha$ denotes the conformally adjusted estimate of the conditional α-quantile.[6] The resulting conformally adjusted intervals accompany the conditional-mean point estimates in subsequent analysis.

C. Emission-Price Shock and Firm Financials

We translate carbon-price increases (0–500 EUR per tCO_2e) into firm financials under short-run, partial-equilibrium assumptions: We abstract from general-equilibrium effects and assume no cost pass-through. Additionally, we abstract from substitution of input factors and technologies, and assume no adjustment of production quantities. The immediate increment in operating costs equals the increase in the carbon price multiplied by the firm's emissions (estimate; using the conditional mean or conformal-adjusted quantiles). We propagate this (operating) cost-shock through the financial statements as follows:

In the profit and loss account, operating costs rise as indicated above. Revenues are held fixed, assuming zero pass-through, so operating profit falls one-for-one with operating costs. Taxes are recomputed from profit before tax using the firm's effective tax rate. If profit before tax is negative, the current tax is set to zero. This yields after-tax profits. In the firm's balance sheet, the change in after-tax profit flows to retained earnings and cash. Current assets (including cash), total assets, and total sources of funds (equity plus liabilities) adjust by the same amount. In the cash-flow statement, operating cash flow changes with after-tax profit, holding trade receivables, inventories, trade payables, and debt constant. In addition to financials, we use financial ratios, covering profitability, interest coverage, solvency, gearing, and profit margins. Financial ratios are recomputed from shock-adjusted numerators and denominators. All mappings and propagation formulas are provided in Table AIII.

[6] While intervals based on raw pinball-loss quantiles are asymptotically valid (conditional and marginal coverage), nominal marginal coverage may deviate in finite samples. In our test set, raw conditional-quantile intervals under-cover by about 4 percentage points. Split conformalisation removes this undercoverage by construction (under the assumption of exchangeability).

D. Default Prediction

We predict a binary firm-default indicator. The default prediction framework uses a European annual firm-level dataset. Variables comprise shock-adjusted financials and derived ratios (which vary between carbon-price shocks) and firm fundamentals (constant between shocks). All variables are used as features in the default-prediction pipelines. Additionally, the dataset includes a binary default label, that is used as target (dependent) variable for the pipelines. A full list of variables is provided in Section E. Analogously to the emission estimation architecture, we construct heterogeneous machine-learning pipelines by systematically combining preprocessing and classification components. Each pipeline follows the same sequence of steps: (i) missing-data indicator, (ii) imputation, (iii) outlier adjustment, (iv) scaling, (v) feature engineering, (vi) feature selection, (vii) resampling for class imbalance, and (viii) classification. At each step, multiple options are available. Again, we generate a large set of pipelines by systematically combining all possible step-option configurations, selecting exactly one option at each step.

For missing data, the options include adding missing-value indicators for numeric features and no inclusion of indicators. Imputation of missing values can be performed using full training sample mean or median replacement. Outlier handling options include no adjustment or winsorisation at 5th/95th percentiles. Scaling is performed with either standard or robust scaler. The next steps include optional application of feature engineering, here squared terms, and applying or bypassing feature selection, here fixed k-best filter, with $k = 25$. The final classification stage includes logistic regression and gradient-boosted trees (XGBoost). Class imbalance is addressed either via class-weighted losses without resampling or via an unweighted loss combined with a hybrid scheme that first oversamples the minority class by generating synthetic samples using k-nearest Neighbors and then removes ambiguous or noisy observations via an edited nearest-neighbour step (applied within the training folds).[7]

[7] Implemented as SMOTE–ENN in *imbalanced-learn* (Lemaitre et al., 2016), i.e., Synthetic Minority Over-

As in the emission estimation, data splitting is organised by firm identifiers so that each firm and all of its annual observations appear in only one partition. For default prediction, firms are randomly allocated into training (70%), calibration (15%) and test (15%) sets. The pipeline selection proceeds again in three stages. First, all candidate pipelines are fitted and evaluated in 15 iterations of Bayesian hyperparameter optimisation (Optuna), using grouped ten-fold cross-validation (grouped by firm) and maximising mean out-of-fold balanced accuracy. This initial optimisation allows us to identify the best pipeline configurations based on the highest balanced accuracy. Second, the best two pipelines are subjected to a more extensive optimisation with 300 iterations of Bayesian hyperparameter search (Optuna) and parameter fitting. In each iteration, hyperparameters are set based on a Bayesian search algorithm, classifier parameters are fitted and pipelines are evaluated as before. Early stopping is applied to terminate trials that underperform running medians of balanced accuracy within or across pipelines. Third, the final hyperparameters per pipeline are set to the values of the iteration with highest balanced accuracy during this second stage, and the corresponding model parameters are re-fitted on the full training set.

Pipelines with class-weighted losses and XGBoost classifier perform best overall.[8] We select the pipeline with the highest cross-validated mean balanced accuracy on the training data and then calibrate the classification threshold for 2022 using the held-out calibration set. Specifically, we choose the threshold that balances precision and recall for the default class, conditional on prior training targeted at balanced accuracy. This yields the final conditional-default classifier model that takes in (if available) firm fundamentals and shock-adjusted financials (of a specific shock value) of a specific year – here 2022 – as well as the financials of the two preceding years, and outputs a binary indicator of firm default (either no default or default) for the selected year. Table II summarises the

sampling Technique (Chawla et al., 2002) followed by Edited Nearest Neighbors (Wilson, 1972).

[8]Pipelines that oversample the minority (default) class with synthetic examples and apply nearest-neighbour editing achieve performance comparable to class-weighted losses but are not applied further due to substantially higher computational cost at similar balanced accuracy. Pipelines with logistic regression perform significantly worse than XGBoost, typically by about 10 percentage points in balanced accuracy on the hold-out test sample.

default model's performance under different specifications and evaluation (sub)samples. Panel a shows strong discrimination of non-defaults but low recall for defaults under the model's default decision rule, implying under-detection of default events. The non-default recall is 0.996 versus 0.140 for defaults, the macro-averaged precision and recall are 0.794 and 0.568, respectively. Overall accuracy is 0.956. Panel b reports performance after threshold calibration to balance precision and recall for the default class. Relative to panel a, default recall increases (with an expected trade-off in precision), i.e., substantially fewer defaulting firms are missed. Within panel b, the evaluation on the 2022 hold-out test subsample assesses generalisation to unseen firms in 2022 (out-of-sample within year), whereas the evaluation on the 2022 full (train, calibration, test) subsample approximates an in-sample–like setting in which shocked firms are more similar to the training population. Accordingly, the unweighted precision and recall (across classes) lie between approximately 0.60 and 0.73 across these two 2022 evaluations; the weighted precision and recall lie above 0.99. Overall, the results indicate that the model is informative for quantifying aggregate impacts, the central objective of portfolio-level climate stress testing.

E. Data

For the emission–estimation framework, we assemble a global firm-level (unbalanced) panel combining firm financials and firm fundamentals with reported firm GHG emissions. Financials and fundamentals are retrieved via Wharton Research Data Services (WRDS) from the Orbis database (Bureau van Dijk, a Moody's Analytics company) between July and September 2025. Reported emissions are obtained through the London Stock Exchange Group (LSEG, formerly Refinitiv) application programming interface in July 2025. The two sources are merged using the Bureau van Dijk identifier (BVDID) and the International Securities Identification Number (ISIN), aligned by fiscal year. The resulting panel comprises 25,100 firm–year observations for 5,647 unique firms over 2010–2023.

Table II. Performance of Default Classification Models

Notes: This table presents the performance of the best-performing default classification model. Panel (a) reports the results for the training using balanced accuracy, that is, the measure of how often the model correctly classifies an observation into default and non-default. Panel (b) reports the results for a balanced precision-recall calibration. The model performance is presented in the form of precision, recall, and support (the number of cases for the respective line). Results are presented for the hold-out subsample, that is, the performance on the part of the sample that the model has not been trained/calibrated on, as well as on the entire sample. The model performance is presented by class (non-default and default) and by average performance (weighted by support and unweighted).

	(a) Balanced accuracy			(b) Balanced precision–recall (year=2022)					
	Hold-out subsample			Hold-out subsample			Full subsample		
Class	Prec.	Recall	Support	Prec.	Recall	Supp.	Prec.	Recall	Supp.
Non-default	0.959	0.996	999,022	0.995	0.996	114,484	0.997	0.996	648,410
Default	0.630	0.140	49,321	0.208	0.200	664	0.435	0.465	3,928
Average	Prec.	Recall	Support	Prec.	Recall	Supp.	Prec.	Recall	Supp.
Unweighted	0.794	0.568	1,048,343	0.602	0.598	115,148	0.716	0.731	652,338
Weighted	0.944	0.956	1,048,343	0.991	0.991	115,148	0.993	0.993	652,338

The dependent variable is reported Scope 1 CO_2-equivalent (CO_2e) emissions in tonnes. The covariate set comprises 76 numeric and 5 categorical variables, plus a firm identifier. Features are selected from the broader Orbis universe based on the share of non-missing observations in the European firm sample used for the default–prediction framework (see below), so as to maximise usable information across firms in the emission–prediction task. Variable definitions are provided in Appendix A (Table AIV); descriptive statistics appear in Table AI.

For the translation of the emission price shock into the financials of the company and for the default–prediction framework, we construct a European[9] firm–level (unbalanced) panel with firm financials, firm fundamentals, and a binary default indicator. The set of features for financials and fundamentals is identical to those used in the emission–estimation framework. We populate these variables with data values retrieved via WRDS from Orbis between July and September 2025 for the European universe. We restrict the analysis

[9] The European sample comprises the EU-27 countries plus Norway, Iceland, Switzerland, the United Kingdom, and Liechtenstein; Albania, Serbia, Montenegro, North Macedonia, Bosnia and Herzegovina, and Kosovo; Ukraine, Belarus, and Moldova; as well as San Marino, Monaco, Andorra, and Vatican City.

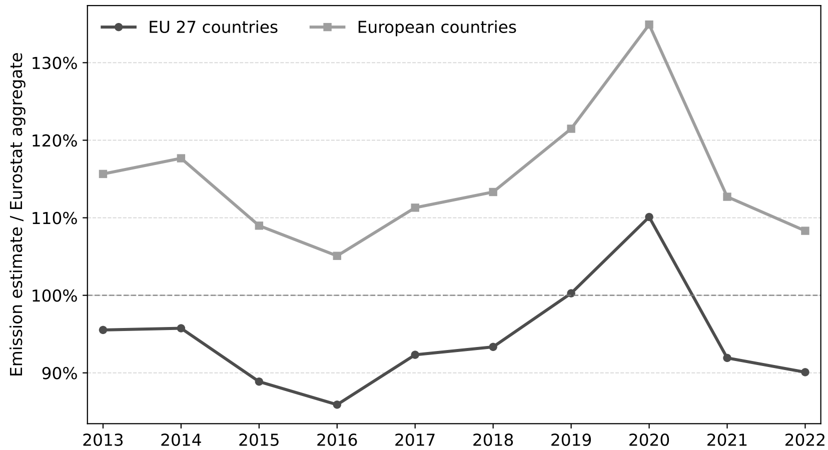

Notes: This figure presents the coverage of aggregate emissions in our predicted sample against the aggregate emissions provided by Eurostat. It presents the coverage for EU27 countries and for all European countries. Note that Eurostat data only covers EU 27 countries.

Figure 4. Coverage of Firm-level Emissions

to firms indebted in 2022, consistent with our stress–testing objective. We augment the panel with firm-level emissions estimates for 2022 obtained from the emission–estimation model with firm financials and fundamentals as input. Based on these inputs, we compute the shock-adjusted financial statement items and derived ratios described in Section C. These calculated variables replace their reported counterparts in the default-model feature set. The final panel spans 2013–2022 and contains 6,484,303 firm–year observations of 1,443,722 firms. To ensure a transparent mapping from emission-price shocks into firm-default prediction, the feature set is restricted to variables whose propagation through accounting identities is unambiguous. Variables potentially affected in multiple and ambiguous ways are excluded. The resulting default–model feature (covariate) set comprises 26 numeric and 3 categorical variables, in addition to a firm identifier. The definitions and formulas of the variables for all all calculated variables, as well as the descriptive statistics for the final default-prediction panel, are provided in Appendix A (Table AIII and Table AII).

To assess the relevance of the firm–level dataset for estimating aggregate effects of emission–price shocks, we benchmark our data along two dimensions against data from Euro-

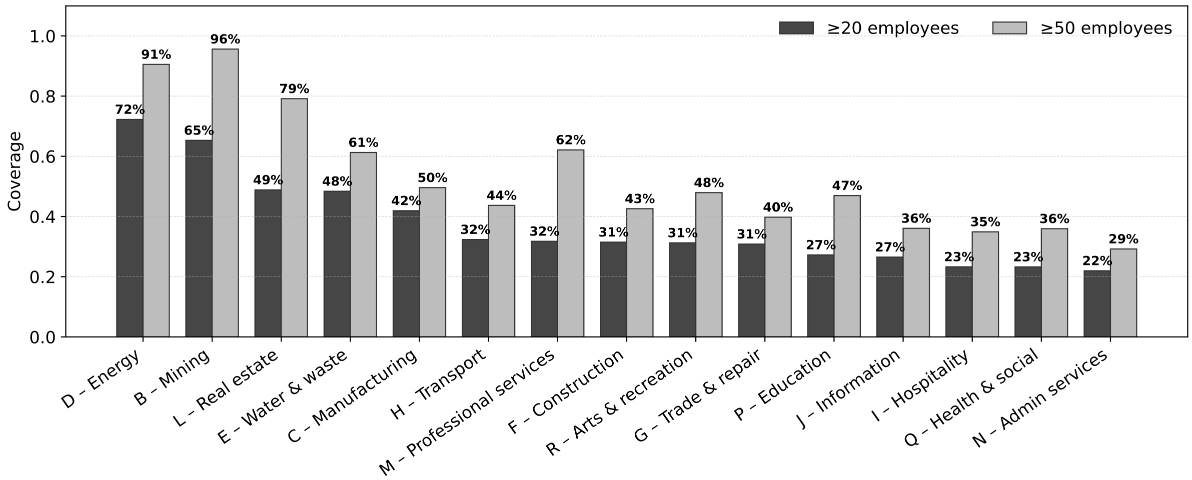

Notes: This figure reports the firm coverage of our sample versus the official statistics by Eurostat, differentiating between firms with ≥ 20 employees and firms with ≥ 50 employees. Reporting is on a sectoral level as defined by NACE Rev. 2.1 level 1. Figures are for the year 2022 and EU27 countries.

Figure 5. Firm Coverage by Sector

stat: (i) coverage of total firm emissions and (ii) coverage of firms.[10] Emissions coverage indicates whether the sample captures the majority of firms' direct cost base under a carbon-price increase, whereas firm coverage indicates the extend to which the impact across the corporate population is represented. It should be noted that Eurostat's emission aggregates are themselves, in part, model-based estimates and should therefore be interpreted accordingly.

At the aggregate EU-27 level, the aggregate estimated emissions closely follow Eurostat during 2013–2022, with emission-coverage ratios typically between 90 and 100% (see Figure 4). Extending the aggregation from the EU-27 to geographic Europe, our sample captures roughly 15% more emissions relative to the EU-27 benchmark (see Figure 4).

At the sectoral level (NACE Rev. 2.1 sections within the EU-27), firm coverage is heterogeneous, yet systematically higher in economically and emissions-relevant sections (see

[10]Eurostat provides harmonised, quality-assured official statistics with consistent time-series coverage and standard NACE classifications across countries, making it a natural reference for aggregate, cross-country, and cross-sector comparisons. We use two coverage notions: (a) emissions coverage (Figure 4) is the ratio of the sum of firm-level emissions in our sample to the corresponding Eurostat aggregate in a given year; (b) firm coverage (Figures 5, 6) is the ratio of the number of firms in our sample to the Eurostat firm count for the respective section or country at a given size threshold. Sources: Eurostat, *Air emissions accounts by NACE Rev. 2 activity* (DOI: 10.2908/ENV_AC_AINAH_R2, retrieved 2025-10-06); Eurostat, *Business demography by size class and NACE Rev. 2 activity* (DOI: 10.2908/SBS_SC_OVW, retrieved 2025-10-07).

Figure 5). In 2022, sections D (Electricity, gas, steam and air conditioning supply), B (Mining and quarrying), and L (Real estate activities) exhibit the highest firm coverage (approximately 80% or higher). A second, larger group—comprising, i.a., E (Water supply; sewerage, waste management and remediation activities), C (Manufacturing), H (Transportation and storage), M (Professional, scientific and technical activities), F (Construction), R (Arts, entertainment and recreation), G (Wholesale and retail trade; repair of motor vehicles and motorcycles), and P (Education)—shows intermediate coverage (about 50%). By contrast, J (Information and communication), I (Accommodation and food service activities), Q (Human health and social work activities), and N (Administrative and support service activities) exhibit lower coverage (around 30%), yet still materially above what is achievable with reported emissions only.[11] Coverage patterns are qualitatively similar when the firm-size threshold is lowered from 50 to 20 employees.

Geographic coverage across European countries is also heterogeneous and reflects cross-country differences in registry and reporting regimes. Figure 6 documents the 2022 cross-section of firm coverage at two size cutoffs (20 and 50 employees). Coverage tends to be higher in northern and southern European countries, while it is lower in France and Germany.

The gap between near-complete emissions coverage at the EU-27 aggregation and more moderate firm coverage within sections is consistent with the distribution of emissions. A small number of large, energy-intensive firms account for a disproportionate share of total emissions,[12]. As a result, an emissions-weighted comparison (aggregate coverage) approaches completeness even if some smaller, low-emission firms are missing. By contrast, the section-level firm-coverage metric is count-weighted: each missing firm contributes equally to the denominator, so gaps among small and medium-sized enterprises depress coverage even when their contribution to total emissions is negligible.

[11] Across all NACE sections, coverage based solely on reported emissions is much lower in our sources: for section B roughly 4%, for D about 2%, and for E about 1%; all other sections are well below 1%.

[12] In our sample, firms with ≥ 250 employees account for approximately 76% of total predicted emissions; firms with 20–249 employees contribute an additional 11%.

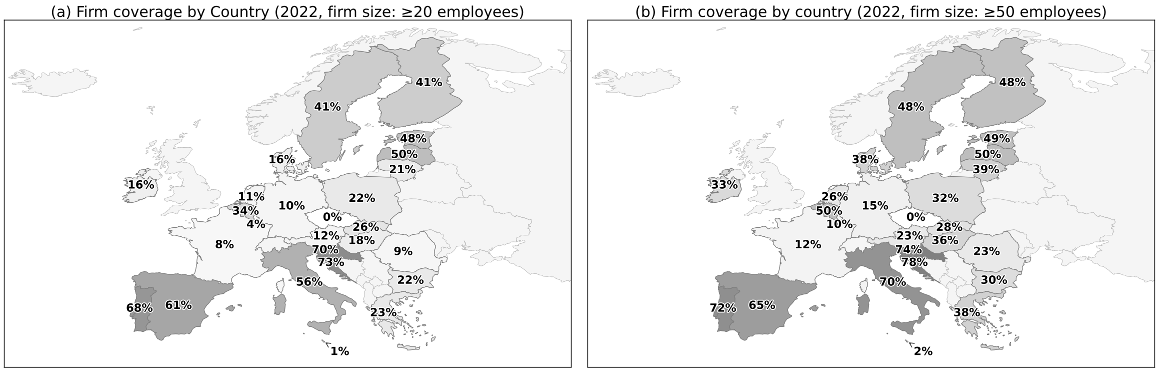

Notes: This figure reports the firm coverage by country of our sample versus the official statistics by Eurostat, differentiating between firms with $\geq$ 20 employees (Panel a) and firms with $\geq$ 50 employees (Panel b). Figures are for the year 2022 and EU27 countries.

Figure 6. Firm Coverage by Country

Overall, the assessment indicates that our firm-level emissions estimates aggregate to near-complete coverage at the EU-27 level, and that firm coverage is highest in emissions-intensive sections while heterogeneous across countries. Importantly, coverage remains markedly higher than that achievable using reported data alone. While not exhaustive, the dataset should allow for a substantially more complete picture of the prospective effects of emission-price shocks on firm default and portfolio level risks.

III. Climate Stress Test Results

A. Aggregate Results

In this section, we present the aggregate effects of a carbon price shock. In the first step, we focus on the impacts on firm financials and profitability based on the methodology presented in Section II.C above, that is, the direct short-run effect of the shock on firms' financials, abstracting from (general-equilibrium) feedback effects and input factor or technology substitution. In the second step, we show how the carbon price shock affects firm defaults, as outlined in Section II.D, and investigate the resulting loan defaults and losses in total assets.

Figure 7 depicts the impact of a climate policy shock on firms' financial viability. As

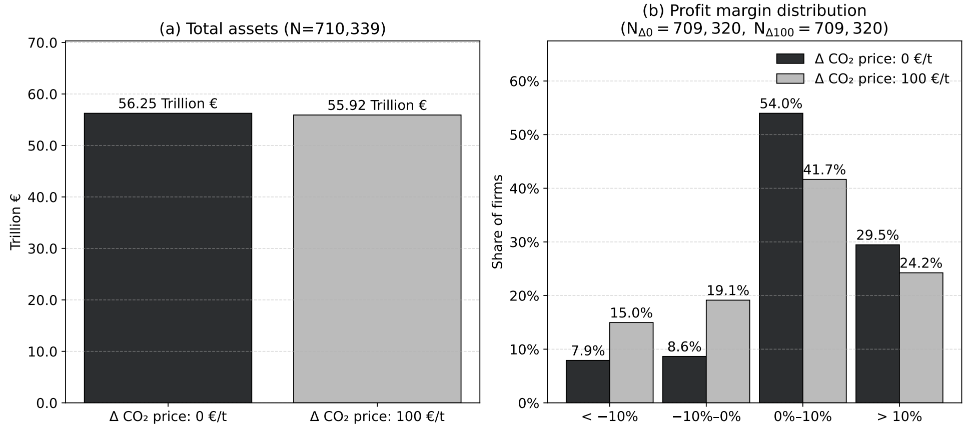

Notes: The figure reports the aggregate effects of a EUR 100 carbon price shock on firm financial viability. Panel (a) reports aggregate total assets and Panel (b) reports the distribution of profit margins for a carbon price shock ($\Delta CO_2\ price = 100\ EUR/t$) and the baseline scenario ($\Delta CO_2\ price = 0\ EUR/t$).

Figure 7. Impact of a Carbon Price Shock on Firm Financial Viability

reported in panel a, the sum of total assets of all firms in the baseline, i.e., $\Delta CO_2\ price = 0\ EUR/t$, is EUR 56.25 trillion, compared to EUR 55.92 trillion after a EUR 100 carbon price shock. This implies that the shock reduces firms' total assets by EUR 330 billion. The effect of the price shock on firm profitability is particularly pronounced. Panel b in Figure 7 depicts the distribution of profit margins, computed as profit before tax over operating revenue. Without the carbon price shock, more than 83% of firms have positive profit margins. A EUR 100 carbon price shock reduces this share to 66%. Consequently, the share of firms incurring losses roughly doubles, from 16.5% to 34%.

Figure 8 explores how these effects differ across firm sizes. With EUR 49.75 trillion of total assets in the baseline, XL companies dominate the aggregate value of total assets (panel a). In contrast, the total assets of large and medium-sized companies amount to EUR 4.74 trillion and EUR 1.77 trillion, respectively (panels c and e). Given this large difference in totals, the shock-induced loss in asset value is relatively similar across firm types. In the case of XL companies, the sum of total assets decreases by EUR 150 billion, compared to EUR 110 billion for large companies and EUR 80 billion for medium-sized companies. This finding indicates that the adverse impacts of a carbon price shock on financial viability decrease with company size. This effect heterogeneity is particularly

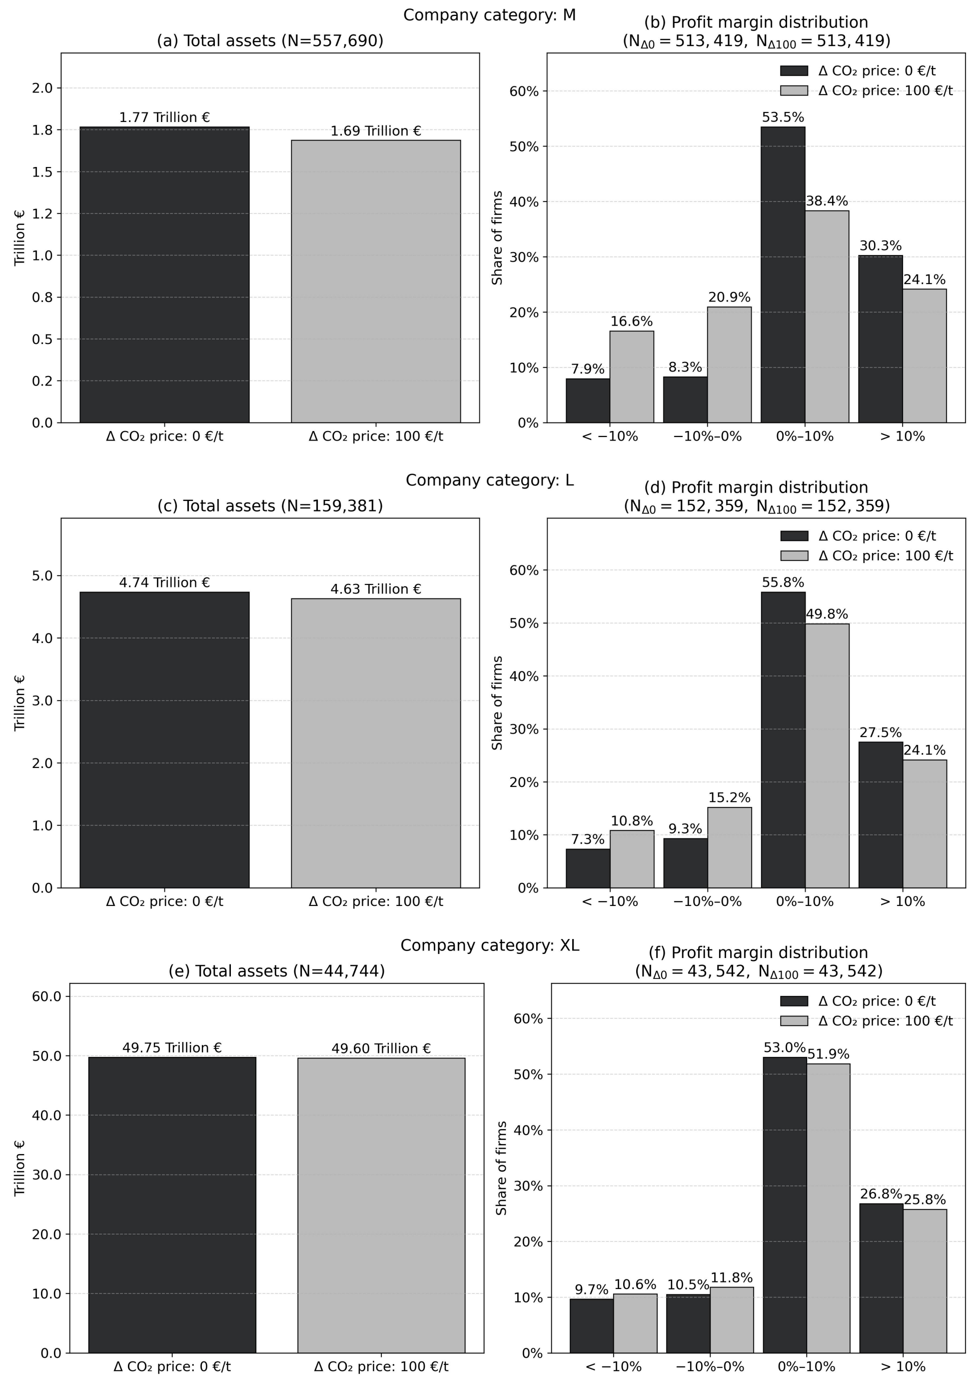

Notes: The figure reports the aggregate effects of a EUR 100 carbon price shock on firm financial viability. Panels (a), (c), and (e) report aggregate total assets and Panels (b), (d), (f) report the distribution of profit margins for a carbon price shock (ΔCO_2 *price* $= 100$ EUR/t) and the baseline scenario (ΔCO_2 *price* $= 0$ EUR/t) differentiating by firm size, defined as: *Very large (XL)*: at least one of revenue $\geq$ 100 m EUR, assets $\geq$ 200 m EUR, employees $\geq 1,000$, or listed; *Large (L)*: at least one of revenue $\geq$ 10 m EUR, assets $\geq$ 20 m EUR, or employees $\geq$ 150, not XL; *Medium (M)*: at least one of revenue $\geq$ 1 m EUR, assets $\geq$ 2 m EUR, or employees $\geq$ 15, not L or XL.

Figure 8. Impact of a Carbon Price Shock on Firm Financial Viability by Firm Size

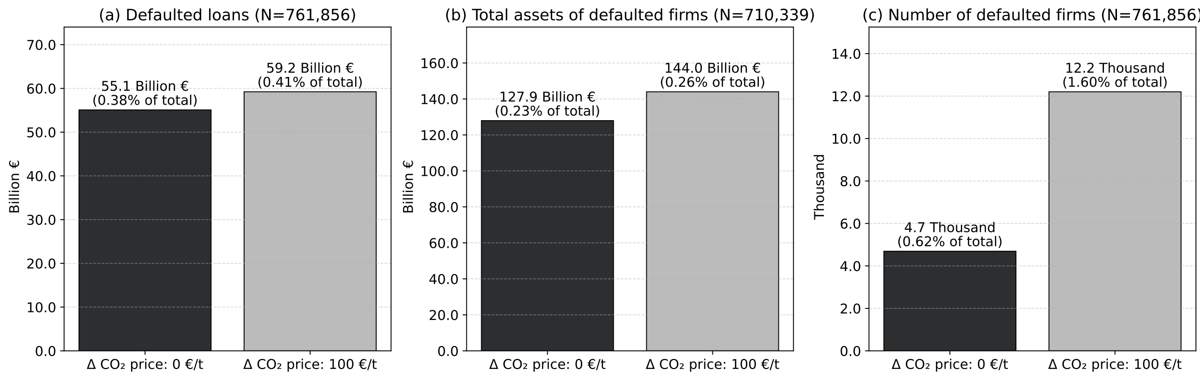

Notes: The figure reports the aggregate effects of a EUR 100 carbon price shock on firm defaults. Panel (a) reports the aggregate defaulted loan volume, Panel (b) reports aggregate total assets of defaulted firms, and Panel (c) reports the number of defaulted firms for a carbon price shock ($\Delta CO_2\ price = 100\ EUR/t$) and the baseline scenario ($\Delta CO_2\ price = 0\ EUR/t$).

Figure 9. Impact of a Carbon Price Shock on Firm Defaults

pronounced in firm profitability. The profit margins of medium-sized firms are affected the most, with the share of firms incurring losses more than doubling from 16% to almost 38% (see panel b). While we observe a weaker but still substantial negative effect on large firms' profitability (panel d), the effects are substantially smaller for XL firms (panel f): the share of profitable (loss-making) firms decreases (increases) by approximately 2 percentage points. These findings indicate that size matters: the financial viability of very large corporations is, on average, only marginally affected, whereas the adverse impacts of carbon price shocks appear particularly strong for smaller firms.

We now go one step further by analysing how the carbon price shock affects firm defaults. Similarly to the analysis above, we compare firm defaults in the baseline and after a EUR 100 carbon price shock. Figure 9 reports the results. As depicted in panel c, the EUR 100 carbon price shock almost triples the number of defaulted firms from 4.7 to 12.2 thousand, which corresponds to 0.6% and 1.6% of all firms[13], respectively. Panel a reports the resulting debt default, computed as the sum of long-term debt of all defaulted firms, considering a full loss given default. The total defaulted loan sum increases by EUR 4.1 billion in the aftermath of a carbon price shock, a 7.9% increase. Finally, panel b reports the total asset loss due to default, computed as the sum of total assets of defaulted firms,

[13]The base year 2022 was characterized by exceptionally low default rates in general.

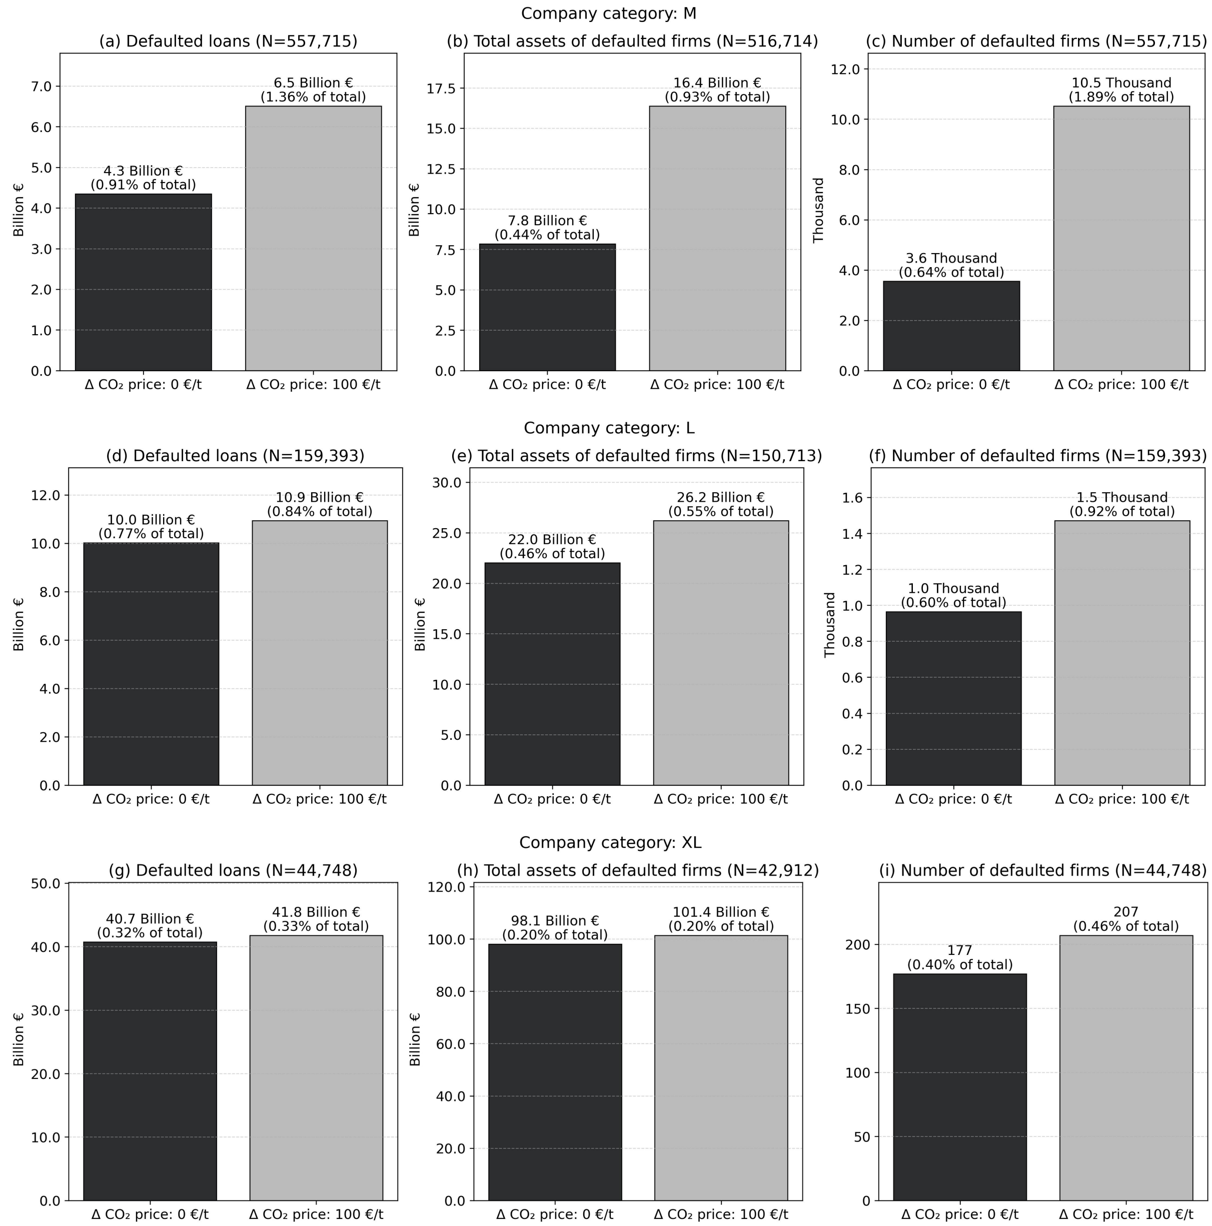

Notes: Notes: The figure reports the aggregate effects of a EUR 100 carbon price shock on firm defaults. Panels (a), (d), and (g) report the aggregate defaulted loan volume, Panels (b), (e), and (h) report aggregate total assets of defaulted firms, and Panels (c), (f), and (i) report the number of defaulted firms for a carbon price shock (ΔCO_2 *price* $= 100$ EUR/t) and the baseline scenario (ΔCO_2 *price* $= 0$ EUR/t) differentiating by firm size, defined as: *Very large (XL)*: at least one of revenue ≥ 100 m EUR, assets ≥ 200 m EUR, employees $\geq 1,000$, or listed; *Large (L)*: at least one of revenue ≥ 10 m EUR, assets ≥ 20 m EUR, or employees ≥ 150, not XL; *Medium (M)*: at least one of revenue ≥ 1 m EUR, assets ≥ 2 m EUR, or employees ≥ 15, not L or XL.

Figure 10. Impact of a Carbon Price Shock on Firm Defaults by Firm Size

which complements the analysis of the number of defaulted firms by providing a proxy for the lost firm value. Compared to the baseline, total assets of defaulted firms are EUR 16 billion higher after the carbon price shock. However, the proportional increase in defaulted debt and loss in assets is notably smaller than the corresponding rise in the number of defaulting firms. This pattern suggests that the additional defaults induced by the carbon price shock are concentrated among relatively smaller firms.

Figure 10 reports firm defaults and the resulting defaulted debt and loss in total assets, confirming this observation. The carbon price shock not only induces the highest increase in firm defaults for medium-sized firms (panel c), both in absolute and relative terms, but also leads to the largest increase in the sum of defaulted debt and total assets of defaulted firms for this firm type. The defaulted debt of medium-sized firms increases by more than EUR 2 billion, compared to around EUR 1 billion for large and extra-large companies, respectively. This pattern is even more pronounced for the total assets of defaulted firms, which rise by more than EUR 16 billion and are thus roughly twice as high as in the baseline for medium-sized companies. This means that not only is the number of firm defaults particularly high for the smallest companies in the sample, but so too is the associated asset value.

B. Sectoral Results

In the next step of the analysis, we investigate potential sectoral differences by differentiating between NACE Rev. 2.1 sections (NACE level 1, A–U) illustrated in Figure 5 above. Specifically, we perform analyses similar to those presented in Section III, but at a sectoral level. Figure 11 depicts the impact of a carbon price shock of EUR 100 on the number of defaulted firms across NACE sections. Panel (a) shows the relative difference, i.e., the number of additional firm defaults in the shock relative to the non-shock scenario. Our results indicate that four sectors experience a particularly high relative increase in defaults. In the mining, trade, transport, and water & waste sectors, the number of defaulted firms increases more than threefold. Although less in magnitude, we still observe a doubling of defaults in, for example, the manufacturing and construction sectors. As depicted in panel b, the absolute number of defaulted firms is particularly large in the

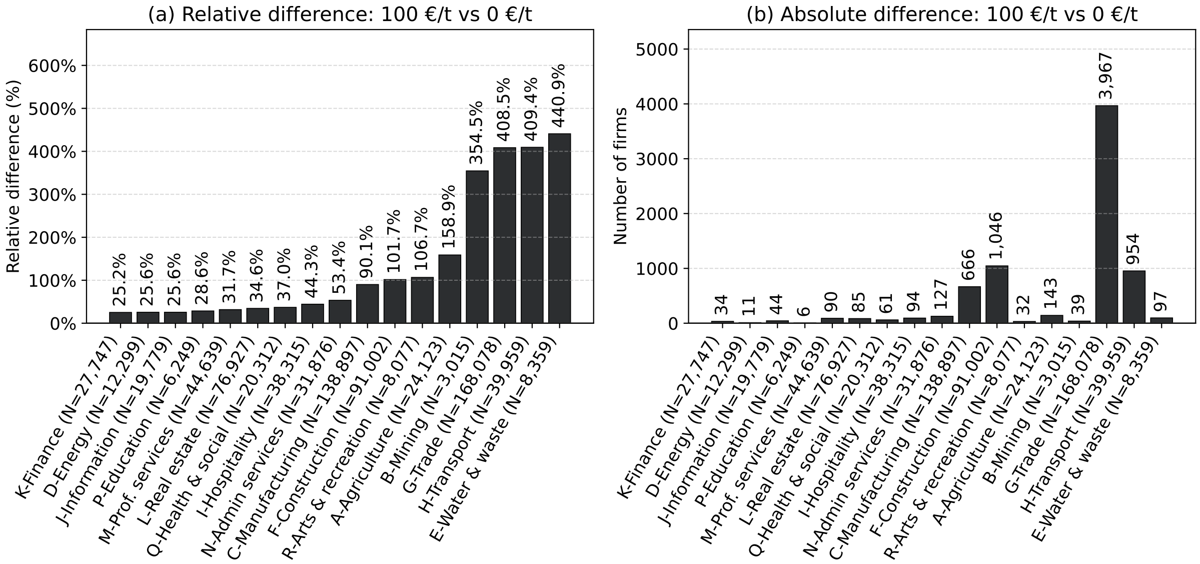

Notes: The figure reports the difference between the number of defaulted firms after a carbon price shock ($\Delta CO_2\ price = 100\ EUR/t$) and in the baseline scenario ($\Delta CO_2\ price = 0\ EUR/t$) across sectors, as defined by NACE Rev. 2.1 level 1. Panel (a) reports the relative difference of defaulted firms. Panel (b) reports the absolute difference of defaulted firms. Sectors are sorted by relative differences.

Figure 11. Impact of a Carbon Price Shock on the Number of Defaulted Firms by Sector

trade sector, where our framework predicts around 4,000 additional defaults due to the carbon price shock, followed by around 1,000 defaulted firms in the construction and manufacturing sectors, respectively.

Figure 12 illustrates the defaulted debt across sectors. Not surprisingly, the three sectors with the highest relative increases in defaulted debt – transport, water & waste, and mining – are also among the top sectors in terms of the number of defaulted firms. However, in these three sectors, the absolute differences in loan volumes are rather small, indicating that defaulting firms in these sectors have relatively low leverage. In absolute terms, we observe the strongest effects in the trade & repair and manufacturing sectors, where loans totalling EUR 1.1 billion and EUR 0.8 billion, respectively, default. Another striking result is the considerable increase in the share of defaulted debt by 21.5% (and EUR 200 billion in absolute terms), even though the relative and absolute numbers of defaulting firms are the second smallest across sectors. This finding indicates that affected energy firms are highly leveraged compared to firms in other sectors.

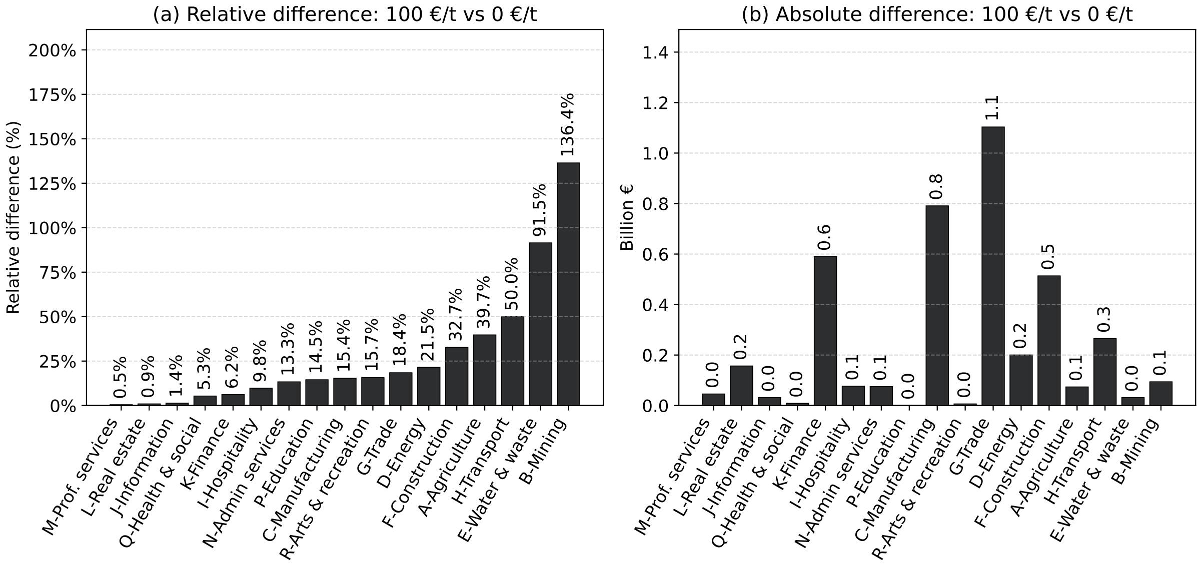

Notes: The figure reports the difference between the volume of defaulted loans after a carbon price shock (ΔCO_2 $price = 100$ EUR/t) and in the baseline scenario (ΔCO_2 $price = 0$ EUR/t) across sectors, as defined by NACE Rev. 2.1 level 1. Panel (a) reports the relative difference of the volume of defaulted loans. Panel (b) reports the absolute difference of the volume of defaulted loans. Sectors are sorted by relative differences. The numbers of firms per NACE section are reported in Figure 11.

Figure 12. Impact of a Carbon Price Shock on Loan Defaults by Sector

Finally, we turn to the asset value of firms defaulting in the aftermath of a carbon price shock. As reported in Figure B4 in the Appendix, the highest relative asset losses due to a carbon price shock can be observed in the transport, agriculture, and mining sectors. However, in absolute terms, the effects in these sectors are rather small. By far the highest reduction in total asset value in the aftermath of a carbon price shock can be observed in the wholesale and resale trade sector (EUR 5.3 billion), followed by manufacturing and construction, with EUR 2.6 billion and EUR 2 billion, respectively. Overall, the sectoral analysis reveals substantial heterogeneity across sectors, with some sectors experiencing particularly high adverse impacts from the carbon price shock compared to the non-shock scenario, such as mining, transport, and trade. Although the relative increase in firm defaults is smaller in the construction and manufacturing sectors, the adverse effects in absolute terms, such as the sum of defaulted debt and the loss of total assets, remain high.[14]

[14] In addition to sectoral differences, we also investigate cross-country differences, which are depicted in Figure B3 in the Appendix.

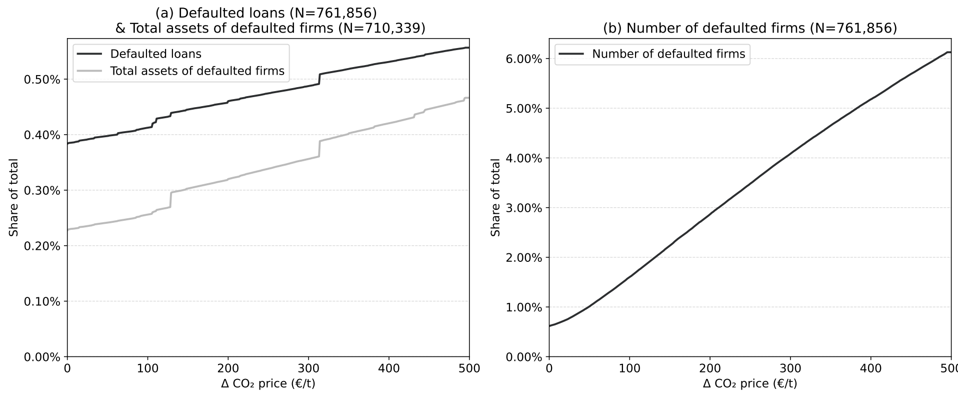

Notes: The figure reports the aggregate effects of varying carbon price shock sizes on the share of defaults. The carbon price shock size is discretely increased by EUR 1 from EUR 0 (baseline) to EUR 500. Panel (a) reports the share of defaulted loan volume and the the share of total assets of defaulted firms for each carbon price shock size. Panel (b) reports the share of the number of defaulted firms for each carbon price shock size.

Figure 13. Defaults at Different Climate Shock Levels

C. Different Carbon Shock Levels

So far, we have considered only a carbon price shock of EUR 100 and compared its impacts on financial viability and default with a baseline without the shock. However, an advantage of our methodology is that we can translate different increases in carbon prices into firm financials and predict the resulting firm-default indicator. Hence, we consider gradual one-EUR carbon price increases from 0–500 EUR per tCO_2e to investigate how the analysed firm outcomes change with the magnitude of the shock.

Figure 13 reports the main results of this analysis. For a carbon price increase of zero, i.e., the baseline, the share of defaulted firms is 0.62% (panel b; see also panel c in Figure 7). After an initially increasing slope, the results reveal an almost linear relationship between the magnitude of the carbon price shock and the share of defaulting firms. For a EUR 500 carbon price shock, our framework predicts the default of more than 6% of firms, which represents a default rate ten times higher than in the baseline. For the share of defaulted debt and the loss of total assets, we also find the effects to increase with shock magnitude. There are, however, two notable differences compared to the share of firms: first, the former increase at a lower rate, indicating that the rise in the share of defaulted firms at higher shock levels is driven by increasing defaults among smaller firms; second,

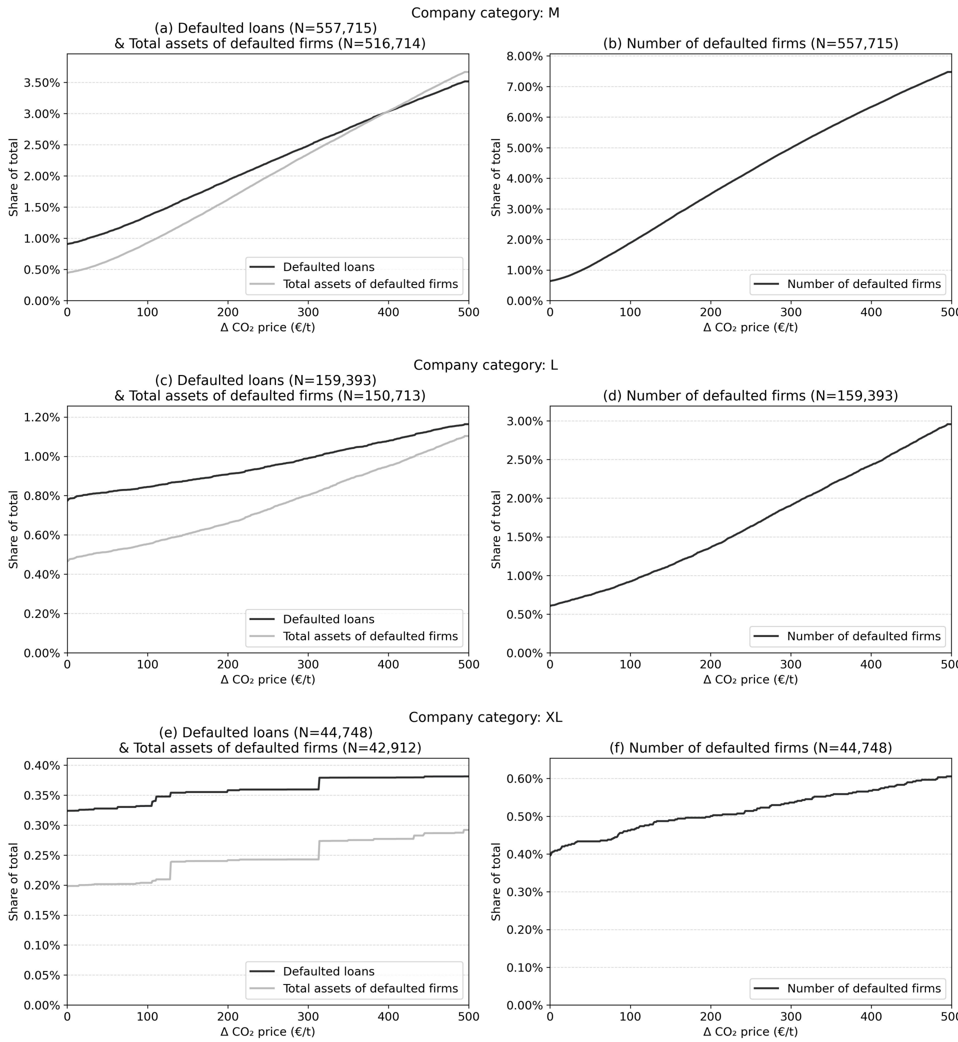

Notes: The figure reports the aggregate effects of varying carbon price shock sizes on the share of defaults. The carbon price shock size is discretely increased by EUR 1 from EUR 0 (baseline) to EUR 500. Panels (a), (c), and (e) report the share of defaulted loan volume and the the share of total assets of defaulted firms for each carbon price shock size. Panels (b), (d), and (f) report the share of the number of defaulted firms for each carbon price shock size differentiating by firm size, which is defined as: *Very large (XL)*: at least one of revenue $\geq$ 100 m EUR, assets $\geq$ 200 m EUR, employees $\geq 1,000$, or listed; *Large (L)*: at least one of revenue $\geq$ 10 m EUR, assets $\geq$ 20 m EUR, or employees $\geq$ 150, not XL; *Medium (M)*: at least one of revenue $\geq$ 1 m EUR, assets $\geq$ 2 m EUR, or employees $\geq$ 15, not L or XL.

Figure 14. Defaults at Different Climate Shock Levels by Firm Size

there are some jumps at different shock magnitudes, driven by very large firms with high outstanding debt switching from non-defaulted to defaulted status.

Figure 14 differentiates the effects by firm size. The results confirm the observations from the aggregate analysis. With respect to the number of defaulted firms, the effect of an increasing carbon price shock magnitude decreases with firm size. For medium-sized firms, an increase in $\Delta CO_2 price$ (EUR/t) from 0 to 500 raises the default rate by around 7 percentage points (panel b). In the case of large companies, this difference in the default rate is relatively smaller, at 2.5 percentage points (panel d). As illustrated in panel f, the share of defaulted XL firms increases from 0.4% to only 0.6%. A similar pattern can be observed for defaulted debt and total assets of defaulted firms. Finally, we observe jumps in the shares of defaulted debt and total assets lost only for XL companies (panel e), which illustrates that the jumps in the aggregate analysis (Figure 13, panel a) are driven by very large firms switching from non-default to default at specific carbon price shocks.

IV. Firm Heterogeneity

An advantage of our framework is that it captures transition risks at the firm level, thereby preserving substantially more information than approaches that rely on more aggregated proxies. In particular, by using heterogeneous firm-level emissions rather than homogeneous sector-level measures, our approach allows for a more accurate assessment of firm-specific exposure to transition risk. This is particularly the case for sectors and industries with pronounced heterogeneity with respect to firm-level emissions. To explore the impact of using firm-level emissions in climate stress testing, we compare our results with two alternative approaches based on division-level emissions. First, we compute the total emissions per NACE division (NACE level 2, 1-99) as the sum of predicted emissions of all firms in that sector. In the first approach, we assign each firm within a division the divisional mean of emissions. In other words, we remove all firm-level heterogeneity at the division level. In the second approach, we assign each firm within a division the revenue-weighted mean of divisional emissions. We then run the default prediction framework and compare the results with those based on our firm-level emission estimates.

Considering the aggregate effects across all firms, we find that the effects of a EUR 100 carbon price shock on firm defaults and the resulting sum of defaulted debt and loss in total assets are substantially larger when using average NACE division emissions as

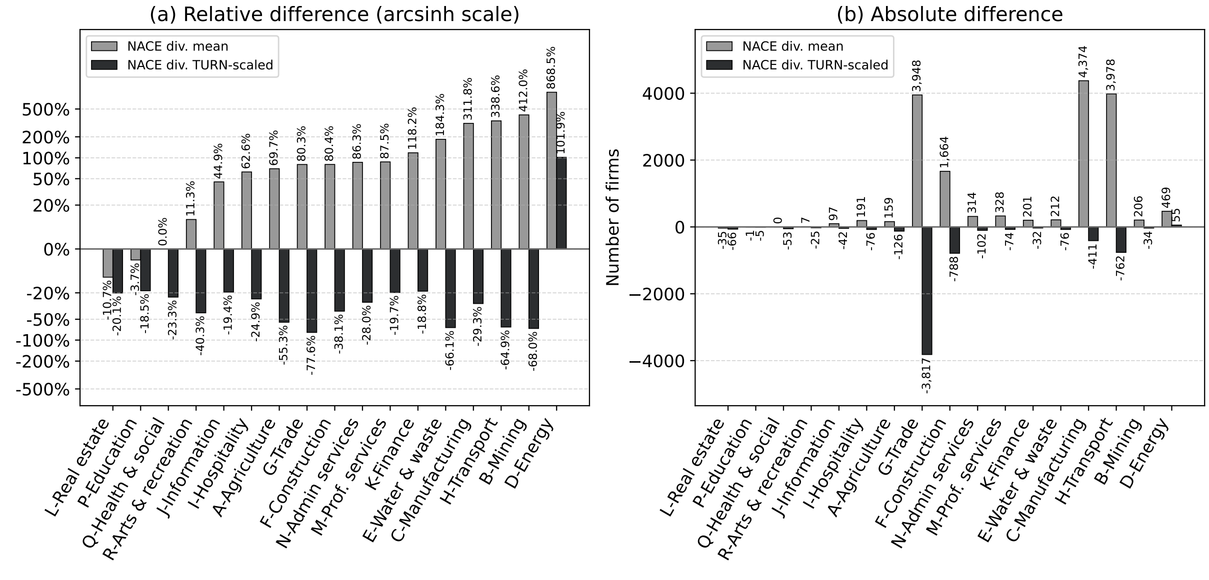

Notes: The figure reports the differences in the number of defaulted firms after a EUR 100 carbon shock using sector-average imputed emissions or revenue-weighted sector-average imputed emissions relative to predicted firm-level emissions. The results are reported at the sectoral level as defined by NACE Rev. 2.1 level 1. Panel (a) reports the relative difference of defaulted firms. Panel (b) reports the absolute difference of defaulted firms. Sectors are sorted by relative differences. The number of firms per NACE section is reported in Figure 11.

Figure 15. Firm Defaults with Predicted Firm Emissions versus Sector Averages

proxies for firm-level emissions (see Figure B5 in the Appendix). This indicates that using sectoral proxies may lead to overestimating the adverse impacts of carbon price shocks. In contrast, using revenue-weighted mean emissions tends to underestimate these impacts. In both cases, the respective over-/underestimation is particularly pronounced for the number of defaulted firms. To further explore the role of firm heterogeneity, we compare the results between NACE sections. The aim of this exercise is to explore which sectors are particularly sensitive to not using firm-level emissions, or in other words, for which sectors the likelihood of over- or underestimating the impacts of shocks is high if (firm-level means of) aggregate emissions, which do not capture firm-level emission heterogeneity, are used as proxies.

Figure 15 depicts the difference in the impact of a carbon price shock of EUR 100 on the number of defaulted firms across NACE sections for both approaches against using firm-level emissions. In line with the aggregate analysis, we find, on average, higher numbers

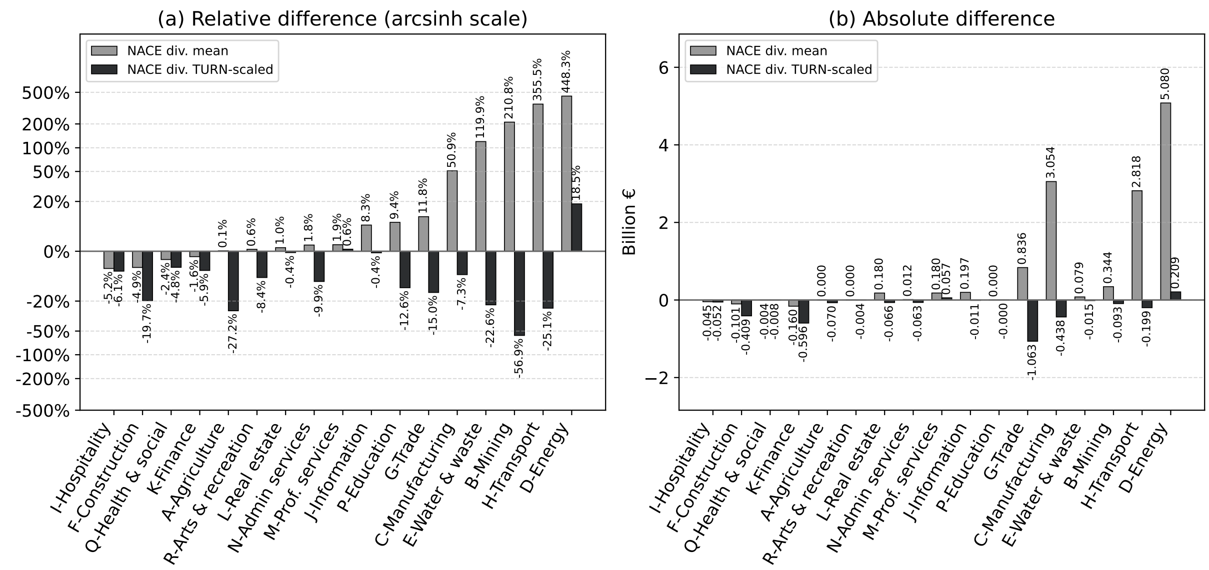

Notes: The figure reports the differences in the volume of defaulted loans after a EUR 100 carbon shock using sector-average imputed emissions or revenue-weighted sector-average imputed emissions relative to predicted firm-level emissions. The results are reported at the sectoral level as defined by NACE Rev. 2.1 level 1. Panel (a) reports the relative difference of the volume of defaulted loans. Panel (b) reports the absolute difference of the volume of defaulted loans. Sectors are sorted by relative differences. The number of firms per NACE section is reported in Figure 11.

Figure 16. Loan Defaults with Predicted Firm Emissions versus Sector Averages

of firm defaults across sectors for divisional mean emissions, and fewer firm defaults for revenue-weighted mean emissions. Relative differences are particularly noteworthy for the energy sector, as the overestimation using mean emissions is particularly high and the use of revenue-weighed emissions also leads to an overestimation in contrast to all other sectors. With respect to the absolute difference, we find the highest (potential) underestimation of firm defaults when using revenue-weighted mean division emissions for the trade sector. The highest overestimation using mean division emissions can be observed for trade, manufacturing, and transport.

Figures 16 and 17 show how the effects of the carbon price shock on defaulted debt and the total assets of defaulted firms differ when using mean division emissions instead of firm-level emissions. An overall pattern is similar to that for defaulted firms; that is, the estimated adverse effects of the shock are larger when using mean divisional emissions for most sectors. Similarly, effects of the shock are smaller for revenue-based mean divisional

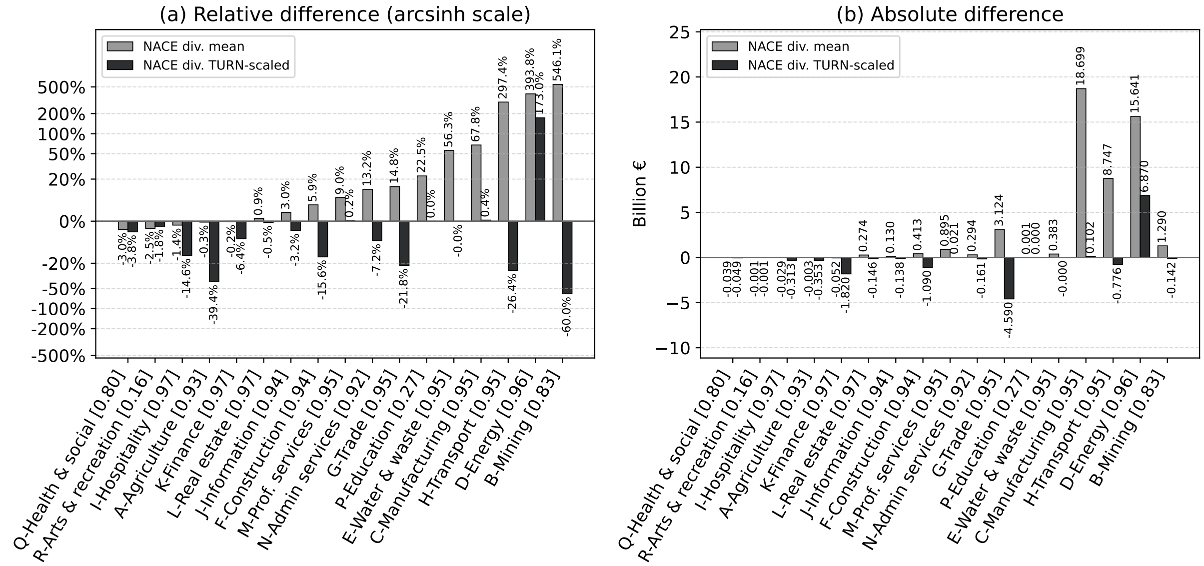

Notes: The figure reports the differences in the total assets of defaulted firms after a EUR 100 carbon shock using sector-average imputed emissions or revenue-weighted sector-average imputed emissions relative to predicted firm-level emissions. The results are reported at the sectoral level as defined by NACE Rev. 2.1 level 1. Panel (a) reports the relative difference of total assets of defaulted firms. Panel (b) reports the absolute difference of total assets of defaulted firms. Sectors are sorted by relative differences. Due to missing data on total assets, the coverage (reported in square brackets) of firm observations as reported in Figure 11, varies by NACE section.

Figure 17. Total Assets of Defaulted Firms with Predicted Firm Emissions versus Sector Averages

emissions. In line with the aggregate analysis (see Figure B5 in the Appendix), the magnitudes of over- and underestimation are higher for the number of firm defaults compared to loan defaults and total assets of defaulted firms.

Overall, this exercise yields two main insights. First, approaches using sectoral means, and hence not capturing firm heterogeneity within a sector, tend to under- or overestimate the adverse effects of a carbon price shock. Second, the magnitude of this potential estimation error, in relative and absolute terms, seems to be highly sector specific. In general, using emission averages based on industry emissions seems to work for sectors with, on average, low emissions and potentially low heterogeneity between firms, such as health & social, whereas firm-level emissions tend to be more important in high-emission sectors with potentially high heterogeneity, such as energy, manufacturing, and transport.

V. Uncertainty

All estimation techniques are subject to model and prediction uncertainty. If large, prediction uncertainties can substantially affect modelling outcomes (Haas et al., 2025) and thus the interpretation of the results. In this section, we investigate the impact of uncertainty in the emission estimation, e.g., due to imperfections of the data or the prediction model, on our main results. We test how those uncertainties affect the default prediction and as such the outcomes of our stress test. As outlined in the methodology (see Section II.B), the uncertainty quantification is based on the 25th and 75th conditional quantile estimates.

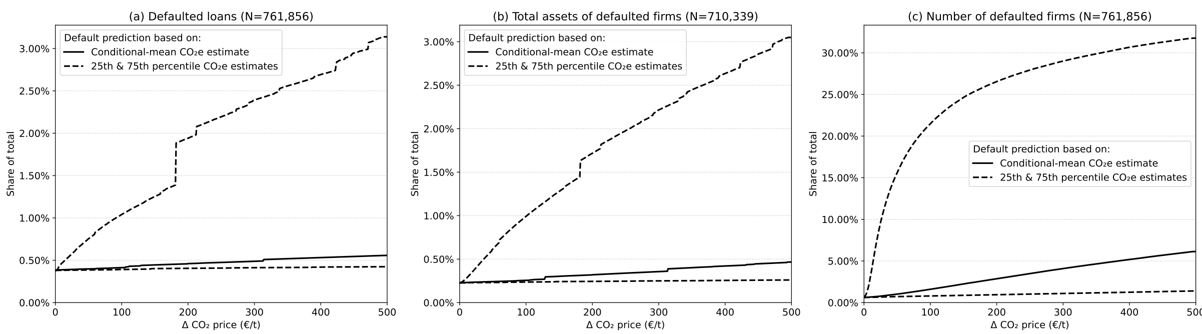

Notes: The figure reports the aggregate effects of varying carbon price shock sizes based on conditional-mean, 25th percentile, and 75th percentile CO_2e estimates. The carbon price shock size is discretely increased by EUR 1 from EUR 0 (baseline) to EUR 500. Panel (a) reports the share of defaulted loan volume, Panel (b) reports the share of total assets of defaulted firms, and Panel (c) reports the share of defaulted firms for each carbon price shock size.

Figure 18. Defaults at Different Climate Shock Levels with Uncertainty

Figure 18 shows the effects of varying carbon price shock sizes based on the conditional-mean CO_2e estimates and – in addition to all previous analyses, such as Figure 13 – the 25th and 75th percentile CO_2e estimates. The 25th and 75th percentiles in the figure are based on perfectly positively correlated prediction errors between firms in the emission estimation. In this case, uncertainty in emission predictions is directly translated into the default prediction. This means that for the 75th percentile emission estimate, the probability of default-related aggregate results exceeding those based on this percentile (upper dashed line in the respective panels) is 25%. Similarly, the probability of the results

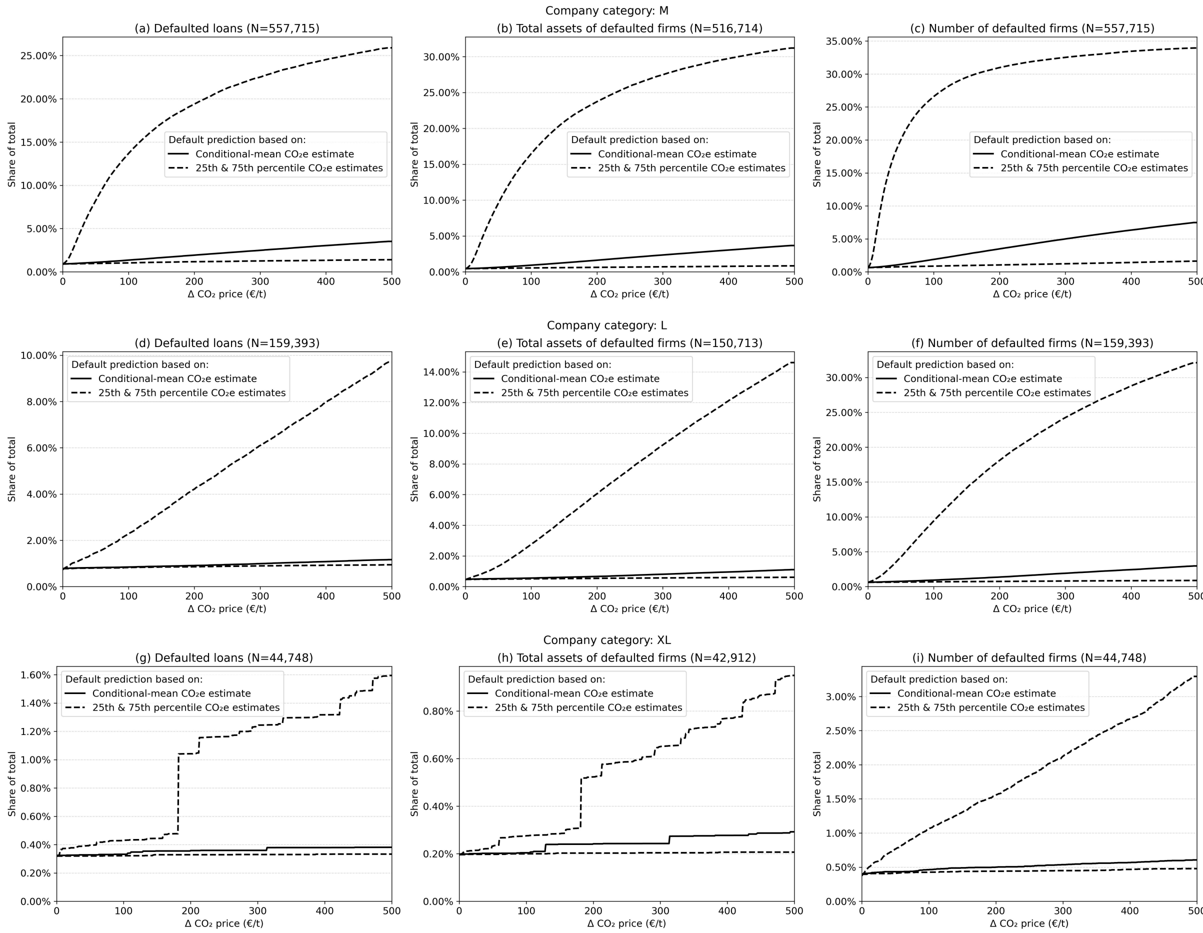

Notes: The figure reports the aggregate effects of varying carbon price shock sizes based on conditional-mean, 25th percentile, and 75th percentile CO_2e estimates. The carbon price shock size is discretely increased by EUR 1 from EUR 0 (baseline) to EUR 500. Panels (a), (d), and (g) report the share of defaulted loan volume, Panels (b), (e), and (h) report the share of total assets of defaulted firms, and Panels (c), (f), and (i) report the share of defaulted firms for each carbon price shock size differentiating by firm size defined as: *Very large (XL)*: at least one of revenue $\geq$ 100 m EUR, assets $\geq$ 200 m EUR, employees $\geq 1,000$, or listed; *Large (L)*: at least one of revenue $\geq$ 10 m EUR, assets $\geq$ 20 m EUR, or employees $\geq$ 150, not XL; *Medium (M)*: at least one of revenue $\geq$ 1 m EUR, assets $\geq$ 2 m EUR, or employees $\geq$ 15, not L or XL.

Figure 19. Defaults at Different Climate Shock Levels with Uncertainty by Firm Size

for the 25th percentile emission estimate falling below those based on this percentile (lower dashed line) is also 25%. If prediction errors are not perfectly positively correlated, then both dashed lines converge towards the conditional mean, due to a reduction in the impact of emission estimation uncertainty on default-related aggregate results.

Data availability and quality affect, indeed, the prediction uncertainty in our stress test. Figure 19 shows that in the case of XL companies, where the availability of CO_2e emission data is the highest, we observe relatively low differences between conditional-mean and quantile estimates. Taking into account the number of defaulted firms after a price shock of EUR 100 (panel i), the share of defaulted firms based on the 75th percentile CO_2e estimate (around 1%) is twice as large as the share based on conditional-mean estimates (around 0.5%). This relationship between estimates is notably stronger for smaller companies with a lower availability of emission data, where the share of firm defaults for L/M companies using the 75th percentile CO_2e estimates is approximately ten/fourteen times higher than the respective predicted share using conditional-mean CO_2e estimates (panels c and f in Figure 19). A similar pattern can be observed for defaulted loans and the total assets of defaulted firms. In summary, this analysis shows that increased emission data availability enhances prediction accuracy, which, in turn, reduces uncertainty in default prediction based on estimated emissions.

This finding has implications for ongoing policy debates. Recent policy developments that scale back or delay corporate climate disclosure requirements, such as the EU's CSRD Omnibus amendments (European Commission, 2025) or the adjustments to the U.S. SEC climate rules (Pinedo et al., 2025), may amplify the uncertainty highlighted above. Reduced reporting, especially among non-listed and medium-sized firms, would further weaken data availability exactly where model uncertainty and default risk are highest. This could further limit the reliability of climate stress tests, as financial institutions and supervisors would have to rely on noisier emission estimates. In this sense, declining disclosure coverage not only increases estimation uncertainty but may, as our analysis suggests, also translate into heightened systemic risk.

VI. Conclusion

This paper proposes machine learning methods and applies these in a micro-prudential climate stress test in the EU. The framework consists of three main steps. First, we estimate Scope 1 GHG emissions from financial and operational data from the BvD Orbis database using a previously trained machine learning model. Second, we analyse how a carbon price shock translates into firms' cost structures and propagates through the profit and loss account, the balance sheet, and cash-flow statements. Third, we build a supervised machine-learning classifier framework that uses the resulting shock-adjusted financials, together with firm fundamentals, to predict binary firm-default events.

Overall, the paper provides valuable insights both on the content level, by shedding light on the firm- and sector-level effects of carbon price shocks, and on the methodological level, by demonstrating how machine learning can enhance climate stress testing and default prediction. Regarding the former, we find that a carbon price shock of EUR 100 substantially weakens the' financial positions of the firms, leading to notable asset losses, reduced profitability, and a nearly doubling of the number of loss-making firms. The shock induces the default of almost 8,000 firms, corresponding to more than EUR 4 billion in additional defaulted loan volume. These adverse effects intensify with the magnitude of the shock and are particularly pronounced among smaller firms, while large corporations exhibit greater resilience. Sectoral differences are considerable, with high-emission industries such as mining and transport being most affected.

The results may be of particular interest when analysing specific firm portfolios. Also banks which are specialized in certain sectors or firm-sizes may want to consider the potentially increased default risk of their portfolio constituents. From a policy perspective, our results suggest that it is important to develop more comprehensive and better quality data infrastructures that cover more than just the largest companies.

On the methodological level, our findings demonstrate that machine learning methods are highly effective tools for climate stress testing. We show that relying on industry-level emission averages tends to understate the true impact of carbon pricing, particularly in emission-intensive sectors characterised by substantial firm-level heterogeneity. In this context, estimating firm-level emissions in the absence of reported data, an issue especially relevant for smaller firms, proves valuable, as it enables the inclusion of SMEs in climate stress tests. However, our analysis of emission estimation uncertainty highlights the importance of suitable prediction models and adequate data availability; otherwise, uncertainty in emission estimates may propagate into the default predictions. As a consequence, ongoing regulatory rollbacks in Europe and the United States on ESG- and climate-related disclosures may undermine the effectiveness of risk management practices in the future, potentially increasing systemic risk.

There are several promising avenues for future research. Future work could examine indirect- and equilibrium effects of carbon price shocks, which are beyond the scope of the present analysis. In addition, the proposed framework could be integrated with bank portfolio data to assess how shock-induced firm defaults might transmit to the financial system and affect overall banking sector stability.

References

Acharya, V.V., Berner, R., Engle, R., Jung, H., Stroebel, J., Zeng, X., Zhao, Y., 2023. Climate Stress Testing. Annual Review of Financial Economics 15. doi:10.1146/annurev-financial-110921-101555.

Akiba, T., Sano, S., Yanase, T., Ohta, T., Koyama, M., 2019. Optuna: A Next-generation Hyperparameter Optimization Framework, in: The 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 2623–2631.

Alogoskoufis, S., Dunz, N., Emambakhsh, T., Hennig, T., Kaijser, M., Kouratzoglou, C., Muñoz, M.A., Parisi, L., Salleo, C., 2021. ECB economy-wide climate stress test Methodology and results. Technical Report. European Central Bank. URL: https://op.europa.eu/en/publication-detail/-/publication/25fa8307-2199-11ec-bd8e-01aa75ed71a1/language-en/format-PDF, doi:10.2866/460490.

Baer, M., Caldecott, B., Kastl, J., Kleinnijenhuis, A.M., Ranger, N., 2022.

TRISK - A Climate Stress Test for Transition Risk. SSRN Electronic Journal URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4254114, doi:10.2139/ssrn.4254114.

Baer, M., Gasparini, M., Lancaster, R., Ranger, N., 2023. Toward a Framework for Assessing and Using Current Climate Risk Scenarios Within Financial Decisions. SSRN Electronic Journal URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4401026, doi:10.2139/ssrn.4401026.

de Bandt, O., Kuntz, L., Pankratz, N., Pegoraro, F., Solheim, H., Sutton, G., Takeyama, A., Xia, F.D., 2025. The effects of climate change-related risks on banks: A literature review. Journal of Economic Surveys 39, 1553–1594. doi:10.1111/joes.12665.

Bank of England, 2022. Results of the 2021 Climate Biennial Exploratory Scenario (CBES). Technical Report. Bank of England.

Battiston, S., Mandel, A., Monasterolo, I., 2019. CLIMAFIN Handbook: Pricing Forward-Looking Climate Risks Under Uncertainty. SSRN Electronic Journal URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3476586, doi:10.2139/ssrn.3476586.

Battiston, S., Mandel, A., Monasterolo, I., Schütze, F., Visentin, G., 2017. A climate stress-test of the financial system. Nature Climate Change 7. doi:10.1038/nclimate3255.

Battiston, S., Martinez-Jaramillo, S., 2018. Financial networks and stress testing: Challenges and new research avenues for systemic risk analysis and financial stability implications. doi:10.1016/j.jfs.2018.03.010.

Battiston, S., Monasterolo, I., Riahi, K., van Ruijven, B.J., 2021. Accounting for finance is key for climate mitigation pathways. Science 372, 918–920. doi:10.1126/science.abf3877.

Bolton, P., Kacperczyk, M., 2023. Global Pricing of Carbon-Transition Risk. The Journal of Finance 78, 3677–3754. doi:10.1111/jofi.13272.

Campiglio, E., Daumas, L., Monnin, P., von Jagow, A., 2023. Climate-related risks in financial assets. Journal of Economic Surveys 37, 950–992. doi:10.1111/joes.12525.

Capasso, G., Gianfrate, G., Spinelli, M., 2020. Climate change and credit risk. Journal of Cleaner Production 266, 121634. doi:10.1016/j.jclepro.2020.121634.

Chawla, N.V., Bowyer, K.W., Hall, L.O., Kegelmeyer, W.P., 2002. SMOTE: Synthetic minority over-sampling technique. Journal of Artificial Intelligence Research 16. doi:10.1613/jair.953.

Chen, T., Guestrin, C., 2016. XGBoost: A scalable tree boosting system, in: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM. pp. 785–794. doi:10.1145/2939672.2939785.

ECB, 2022. 2022 climate risk stress test. Technical Report. European Central Bank. Frankfurt am Main. URL: https://www.bankingsupervision.europa.eu/ecb/pub/pdf/ssm.climate_stress_test_report.20220708~2e3cc0999f.en.pdf, doi:10.2866/97350.

Emambakhsh, T., Fuchs, M., Kordel, S., Kouratzoglou, C., Lelli, C., Pizzeghello, R., Salleo, C., Spaggiari, M., 2023. The Road to Paris: Stress Testing the Transition Towards a Net-Zero Economy. SSRN Electronic Journal doi:10.2139/ssrn.4564374.

European Commission, 2025. Proposal for a Directive of the European Parliament and of the Council amending Directives 2006/43/EC, 2013/34/EU, (EU) 2022/2464 and (EU) 2024/1760 asregards certain corporate sustainability reporting and due diligence requirements.

Grippa, P., Mann, S., 2021. Climate-Related Stress Testing: Transition Risks in Norway. SSRN Electronic Journal doi:10.2139/ssrn.3758065.

Haas, C., Moslener, U., Rink, S., 2025. Corporate Sustainability Data, Machine Learning, and Prediction Uncertainty. Unpublished Working Paper .

Hoerl, A.E., Kennard, R.W., 1970. Ridge Regression: Biased Estimation for Nonorthogonal Problems. Technometrics 12. doi:10.1080/00401706.1970.10488634.

Jung, H., Engle, R., Berner, R., 2021. CRISK: Measuring the Climate Risk Exposure of the Financial System. FRB of New York Staff Report 977. doi:10.2139/ssrn.3931516.

Kainth, D., Melin, L., Rebonato, R., 2024. Climate Scenario Analysis and Stress Testing for Investors: A Probabilistic Approach. An EDHEC-Risk Climate Impact Institute Publication .

Kempa, K., Moslener, U., 2022. Climate and Environmental Policy Risk and Debt. SSRN Electronic Journal doi:10.2139/ssrn.4274711.

Koberle, A., Ganguly, G., Ostrovnaya, A., 2021. A guide to building climate-financial scenarios for financial institutions. Technical Report. Imperial College London. London.

Lemaitre, G., Nogueira, F., Aridas, C.K., 2016. Imbalanced-learn: A Python Toolbox to Tackle the Curse of Imbalanced Datasets in Machine Learning. Journal of Machine Learning Research 7, 1–5.

Monasterolo, I., 2020. Climate change and the financial system. Annual Review of Resource Economics 12. doi:10.1146/annurev-resource-110119-031134.

Nguyen, Q., Diaz-Rainey, I., Kuruppuarachchi, D., McCarten, M., Tan, E.K., 2023. Climate transition risk in U.S. loan portfolios: Are all banks the same? International Review of Financial Analysis 85. doi:10.1016/j.irfa.2022.102401.

Nieto, M.J., 2019. Banks, climate risk and financial stability. Journal of Financial Regulation and Compliance 27. doi:10.1108/JFRC-03-2018-0043.

Nordhaus, W., 2018. Projections and Uncertainties about Climate Change in an Era of

Minimal Climate Policies. American Economic Journal: Economic Policy 10, 333–360. doi:10.1257/pol.20170046.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: Machine learning in Python. Journal of Machine Learning Research 12.

Pinedo, A., Walsh, L., Juarez, C., 2025. Regulatory Climate Shift: Updates on the SEC Climate-Related Disclosure Rules. Technical Report. Harvard Law School Forum on Corporate Governance.

Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A.V., Gulin, A., 2018. Catboost: Unbiased boosting with categorical features, in: Advances in Neural Information Processing Systems, pp. 1–23.

Reinders, H.J., Schoenmaker, D., van Dijk, M., 2023. A finance approach to climate stress testing. Journal of International Money and Finance 131. doi:10.1016/j.jimonfin.2022.102797.

Reinders, H.J., Schoenmaker, D., van Dijk, M., 2025. Climate risk stress testing: A critical survey and classification. Journal of Climate Finance 10, 100061. doi:10.1016/j.jclimf.2025.100061.

Romano, Y., Patterson, E., Candès, E.J., 2019. Conformalized quantile regression, in: Advances in Neural Information Processing Systems, pp. 1–19.

Seltzer, L., Starks, L.T., Zhu, Q., 2022. Climate Regulatory Risk and Corporate Bonds. SSRN Electronic Journal doi:10.2139/ssrn.4098316.

Tol, R.S., 2023. Social cost of carbon estimates have increased over time. Nature Climate Change 13, 532–536. doi:10.1038/s41558-023-01680-x.

Vermeulen, R., Schets, E., Lohuis, M., Kölbl, B., Jansen, D.J., Heeringa, W., 2019.

The Heat Is on: A Framework for Measuring Financial Stress Under Disruptive Energy Transition Scenarios. SSRN Electronic Journal doi:10.2139/ssrn.3346466.

Wilson, D.L., 1972. Asymptotic Properties of Nearest Neighbor Rules Using Edited Data. IEEE Transactions on Systems, Man and Cybernetics 2. doi:10.1109/TSMC.1972.4309137.

Xu, H.C., Li, T.M., Dai, P.F., Nguyen, D.K., Zhou, W.X., 2024. Stress testing climate risk: A network-based analysis of the Chinese banking system. Journal of International Money and Finance 149, 103207. doi:10.1016/j.jimonfin.2024.103207.

ONLINE APPENDIX

A. Variables & Descriptive Statistics

Table AI. Descriptive Statistics of Input Variables for Emission Prediction Model

Numeric variables	N	Mean	Std	Min	25%	50%	75%	Max
Capital and Reserves (bn USD)	24,884	0.98	4.77	-1.25	0.01	0.14	0.66	200.99
Cash and Cash Equivalents (bn USD)	25,035	1.65	10.60	-0.00	0.09	0.37	1.25	1465.60
Cash Flow (bn USD)	21,490	1.53	9.47	-19.05	0.08	0.31	1.00	1196.24
Cash Flow from Operations (USD)	21,104	15.25	16.74	-99.73	6.52	12.14	21.61	99.94
CO2 Equiv. Emissions Scope 1 (Mt CO2e)	25,100	2.98	14.75	0.00	0.01	0.06	0.48	764.93
Cost of Sales (bn USD)	22,540	7.26	21.24	-2.76	0.37	1.54	5.65	452.78
Creditor Days (days)	24,864	43.08	44.49	0.00	20.00	34.00	53.00	999.00
Creditors (bn USD)	24,926	1.28	3.56	-0.14	0.06	0.24	0.94	95.14
Current Assets (bn USD)	25,100	5.64	24.95	-0.00	0.44	1.50	4.54	3216.60
Current Liabilities (bn USD)	25,100	4.44	15.31	0.00	0.29	1.01	3.37	1537.91
Current Ratio (k %)	25,091	0.18	0.18	0.00	0.10	0.14	0.20	7.29
Debtor Days (days)	25,000	54.59	50.35	0.00	26.00	47.00	70.00	995.00
Debtors (bn USD)	25,051	1.55	11.29	-0.00	0.08	0.34	1.16	945.67
Depreciation (bn USD)	21,753	0.73	2.40	-3.79	0.03	0.13	0.47	113.30
Earnings Before Tax and Amortization (bn USD)	23,128	1.86	10.67	-25.51	0.12	0.41	1.29	1349.45
Earnings to Turnover (%)	22,883	19.42	18.58	-99.95	8.39	15.26	26.59	99.82
EBITDA Margin (%)	24,561	12.23	17.20	-99.69	4.57	9.65	17.65	99.96
Extraordinary Items (mn %)	20,125	-7.27	779.77	-13700.00	-11.34	0.00	0.02	80066.15
Finance Expenses (bn USD)	24,018	0.21	1.14	-24.50	0.01	0.04	0.18	76.88
Finance Profit/Loss (mn USD)	24,336	-90.69	1155.40	-76803.70	-112.95	-19.81	0.46	81360.15
Finance Revenue (bn USD)	21,918	0.12	1.09	-3.97	0.00	0.01	0.05	100.41
Fixed Assets (bn USD)	25,100	11.23	30.42	0.00	0.75	2.64	9.34	1772.25
Gearing Ratio (k %)	24,269	0.11	0.12	0.00	0.04	0.08	0.13	1.00
Gross Margin (%)	22,739	43.12	24.85	-87.78	23.89	38.64	59.45	100.00
Gross Profit (bn USD)	22,840	4.17	35.98	-23.39	0.32	1.03	3.10	5190.04
Intangible Fixed Assets (bn USD)	25,075	2.90	11.49	-0.71	0.04	0.28	1.55	796.50
Interest Cover (times)	22,289	23.39	72.13	-97.83	2.45	6.18	16.19	998.42
Interest Expenses (bn USD)	22,637	0.19	0.49	-0.27	0.01	0.04	0.15	15.45
Liquid Ratio (k %)	25,068	0.14	0.17	0.00	0.07	0.11	0.15	8.76
Loans (bn USD)	24,914	0.91	3.64	-7.52	0.01	0.11	0.59	136.02
Long-Term Debt (bn USD)	25,100	3.81	9.86	-0.24	0.14	0.77	3.10	191.00
Material Expenses (bn USD)	4,552	5.02	15.93	-0.01	0.14	0.68	3.38	298.35
Net Asset Turnover (times)	25,045	1.31	2.43	0.00	0.52	0.95	1.56	301.80

Continued on next page

Table AI. Descriptive Statistics of Input Variables for Emission Prediction Model

Numeric variables	N	Mean	Std	Min	25%	50%	75%	Max
Net Current Assets (bn USD)	21,713	1.13	12.55	-40.32	0.01	0.23	0.98	1678.70
Non-Current Liabilities (bn USD)	25,098	5.93	15.97	-3.30	0.27	1.20	4.54	371.18
Number of Employees (k count)	18,799	30.85	75.34	0.00	2.80	9.48	29.01	2300.00
Operating Profit/Loss (bn USD)	24,914	1.15	9.00	-29.01	0.06	0.25	0.83	1236.15
Operating Revenue (bn USD)	25,100	11.07	44.46	-2.17	0.85	2.80	9.14	5361.27
Other Current Assets (bn USD)	25,094	3.29	23.15	-1.76	0.18	0.65	2.08	2250.58
Other Current Liabilities (bn USD)	24,925	2.25	10.63	0.00	0.11	0.44	1.53	1323.92
Other Fixed Assets (bn USD)	25,087	2.91	12.12	-1.32	0.07	0.36	1.66	533.52
Other Non-Current Liabilities (bn USD)	25,100	2.12	7.12	-4.06	0.06	0.30	1.22	180.18
Other Operating Expenses (bn USD)	24,047	3.12	26.96	-5.75	0.21	0.70	2.22	3953.89
Other Shareholders' Funds (bn USD)	24,884	5.50	25.61	-114.56	0.33	1.32	4.38	3239.49
Profit Margin (%)	24,528	10.17	17.26	-99.97	3.38	8.31	15.77	99.91
Profit/Loss After Tax (bn %)	25,017	0.80	8.58	-23.28	0.03	0.15	0.56	1082.94
Profit/Loss Before Tax (bn USD)	25,089	1.08	10.27	-28.88	0.04	0.20	0.73	1317.51
Profit/Loss for the Financial Year (bn %)	25,061	0.76	7.55	-22.44	0.03	0.15	0.54	1082.94
Property, Plant and Equipment (PPE) (k USD)	18,119	0.14	0.55	-0.10	0.01	0.03	0.07	9.61
Provisions (bn USD)	18,665	1.01	3.76	-0.04	0.01	0.09	0.47	101.06
Return on Assets (%)	25,034	4.38	8.26	-97.50	1.66	4.18	7.47	89.78
Return on Capital Employed (at) (%)	22,603	7.99	18.03	-929.67	4.11	7.53	12.33	647.59
Return on Capital Employed (pt) (%)	22,600	10.47	20.23	-933.65	5.27	9.54	15.66	836.46
Return on Equity (%)	24,634	10.55	40.16	-968.96	4.32	9.96	17.05	872.84
Return on Shareholders' Funds (%)	24,656	14.90	43.63	-911.82	6.37	13.42	22.30	994.77
Return on Total Assets (%)	25,061	6.01	9.37	-98.14	2.49	5.65	9.70	95.05
Share Capital (USD)	17,124	15.59	14.04	0.00	5.07	11.86	21.83	100.00
Shareholders' Funds (bn USD)	25,099	6.48	26.25	-18.32	0.55	1.82	5.41	3262.79
Shareholders' Funds to Long-Term Liab. (times)	15,538	54.09	24.47	0.03	35.41	53.31	73.66	100.00
Shareholders' Funds to PPE (k times)	17,859	0.54	1.12	0.00	0.09	0.20	0.44	9.99
Shareholders' Funds to Total Assets (%)	25,076	5.06	23.71	-53.95	0.82	1.55	3.24	983.09
Solvency Ratio (%)	25,072	43.48	19.51	-99.28	30.90	43.44	56.37	99.33
Staff Expenses (bn USD)	17,258	1.43	25.82	-9.68	0.06	0.24	0.80	2480.54
Stock Turnover (times)	22,472	34.14	84.38	0.00	5.68	9.20	23.23	996.53
Stocks (bn USD)	25,076	1.13	3.28	-0.18	0.03	0.21	0.89	81.72
Tangible Fixed Assets (bn USD)	25,098	5.42	15.41	0.00	0.23	0.97	3.97	442.23
Taxation (bn %)	24,893	0.27	2.16	-27.60	0.01	0.05	0.17	234.56
Total Assets (bn USD)	25,100	16.87	50.96	0.00	1.44	4.59	14.41	4988.86
Total Assets to PPE (k times)	18,629	2.25	7.34	0.00	0.25	0.49	1.18	99.18
Total PPE (k USD)	18,708	0.92	3.54	0.00	0.21	0.35	0.65	94.69
Total Shareholders' Funds and Liab. (bn USD)	25,099	16.87	50.96	0.00	1.44	4.59	14.41	4988.86
Turnover (bn USD)	22,706	11.27	45.77	-1.99	0.87	2.88	9.47	5348.96
Working Capital (bn USD)	21,910	1.36	7.22	-87.99	0.06	0.32	1.17	888.06

Continued on next page

Table AI. Descriptive Statistics of Input Variables for Emission Prediction Model

Numeric variables	N	Mean	Std	Min	25%	50%	75%	Max
Working Capital to PPE (k times)	16,037	0.11	0.36	-0.10	0.02	0.05	0.10	9.21
Year	25,100	2017	4	2010	2014	2018	2020	2023
Year of Incorporation	19,585	1981	32	1792	1968	1992	2003	2023

Categorial variables	N	Number of unique values
Category of Company	25,100	3
Country (ISO Code)	25,100	74
Legal Form	25,069	7
Listed on Stock Exchange	25,100	2
NACE Code (2-digit)	25,077	84

Table AII. Descriptive Statistics of Input Variables for Default Prediction Model

Numeric variables	N	Mean	Std	Min	25%	50%	75%	Max
Cash c. (mn EUR)	6,484,303	4.94	167.70	-61391.12	0.03	0.16	0.65	75746.04
Cash Flow c. (bn %)	6,067,297	0.43	16.12	-6139.21	0.00	0.02	0.06	6140.30
Cost of Goods Sold c. (mn EUR)	6,779,151	46.30	1178.72	0.00	0.89	2.19	7.49	466147.71
Current Assets (mn EUR)	6,994,649	27.24	1081.58	0.00	0.51	1.40	4.57	905001.69
Current Liabilities (mn EUR)	6,993,656	21.34	898.10	0.00	0.30	0.87	3.00	577480.06
Current Ratio (k %)	6,882,532	0.30	0.69	0.00	0.10	0.15	0.24	10.00
Delta Profit/Loss After Tax (k EUR)	6,692,975	-87.23	31069.75	-61392724.60	-0.00	0.00	0.00	10559000.00
Fixed Assets (mn EUR)	6,993,101	47.15	1662.39	0.00	0.24	1.12	4.62	1148074.10
Gearing Ratio c. (k %)	6,196,008	10.55	4638.40	0.00	0.01	0.04	0.12	6588353.98
Interest Coverage c. (k times)	6,233,184	0.27	61.19	-22049.27	0.00	0.00	0.02	80459.25
Liquid Ratio (k %)	6,871,978	0.24	0.62	0.00	0.06	0.11	0.18	10.00
Net Current Assets (mn EUR)	4,320,579	8.37	548.60	-147826.30	-0.00	0.39	1.82	862024.36
Non-Current Liabilities (mn EUR)	6,993,728	26.28	763.72	0.00	0.14	0.56	2.42	249004.22
Operating Profit/Loss c. (mn EUR)	6,779,151	3.07	158.89	-140029.30	0.01	0.10	0.40	162062.77
Profit Margin c. (k %)	6,691,449	-1.99	445.12	-790523.91	0.00	0.00	0.01	0.10
Profit/Loss After Tax c. (mn EUR)	6,772,651	1.81	144.48	-143853.32	0.00	0.06	0.28	158098.59
Profit/Loss Before Tax c. (mn EUR)	6,772,651	2.51	163.34	-143853.32	0.01	0.08	0.36	158098.59
Return Shareholders' Funds c. (k %)	6,196,008	0.50	698.16	-529822.45	0.00	0.01	0.04	1261270.33
Return Total Assets c. (%)	6,690,896	7.24	11717.98	-18172080.33	0.74	3.91	9.29	17414190.00
Shareholders' Funds c. (mn EUR)	6,655,968	24.62	1261.46	-61359.71	0.22	0.90	3.42	1148079.23
Solvency Ratio c. (%)	6,194,710	46.78	19041.48	0.00	18.23	35.18	55.82	47319630.00
Taxation c. (mn EUR)	6,772,651	0.69	34.90	0.00	0.00	0.01	0.06	39302.00

Continued on next page

Table AII. Descriptive Statistics of Input Variables for Default Prediction Model

Numeric variables	N	Mean	Std	Min	25%	50%	75%	Max
Total Assets c. (mn EUR)	6,691,315	70.92	2156.67	-61356.21	1.27	3.35	10.33	1148079.34
Total Shareholders' Funds c. (mn EUR)	6,654,908	67.83	2138.52	-61356.21	1.26	3.33	10.16	1148079.34
Working Capital (mn EUR)	6,257,224	6.28	298.55	-276663.40	0.03	0.33	1.47	325706.06
Year	6,994,751	2018	3	2013	2015	2018	2020	2022
Categorial variables	**N**	**Number of unique values**						
Category of Company	6,994,751	3						
Country ISO Code	6,994,751	44						
NACE 2-Digit Code	6,965,748	89						

Table AIII. Definitions of Variables Used in the Default Prediction Model

Variable name	Code	Description	Formula (if calculated)
Cash c.	CASH_calc	Cash and cash equivalents.	$\mathtt{CASH\_calc} = \mathtt{CASH} + \Delta\mathtt{PLAT}$
Cash Flow c.	CF_calc	Cash flow from operations.	$\mathtt{CF\_calc} = \mathtt{PLAT\_calc} + \mathtt{DEPR}$
Category of Company		Company size class (SM, L, XL).	–
Cost of Goods Sold c.	COST_shock	Cost of goods sold.	$\mathtt{COST\_shock} = \mathtt{COST\_calc} + (\mathtt{CO2e\ estimation} \times \Delta\mathtt{emission\_price})$; with $\mathtt{COST\_calc} = (\mathtt{OPRE} - \mathtt{OPPL})_{\geq 0}$
Country ISO Code		Country of incorporation (ISO code).	–
Current Assets	CUAS	Total current assets.	–
Current Liabilities	CULI	Total current liabilities.	–
Current Ratio	CURR	Current assets / current liabilities .	–
Default Indicator	`target`	Dummy (1 = default, 0 = no default).	–
Delta Profit/Loss After Tax	Delta_PLAT	Change in profit after tax (shock vs. baseline).	$\Delta\mathtt{PLAT} = \mathtt{PLAT\_calc} - \mathtt{PLAT}$
Fixed Assets	FIAS	Net of accumulated depreciation.	–
Gearing Ratio c.	GEAR_calc	Long-term debt / shareholders' funds .	$\mathtt{GEAR\_calc} = \frac{\mathtt{LTDB}}{\mathtt{SHFD\_adj}} \times 100$
Interest Coverage c.	IC_calc	Operating profit / financial expenses.	$\mathtt{IC\_calc} = \frac{\mathtt{OPPL\_calc}}{\mathtt{FIEX}}$
Liquid Ratio	LIQR	(Current assets – inventories) / current liabilities.	–
NACE 2-Digit Code		NACE Rev. 2, two-digit.	–
Net Current Assets	NCAS	Current assets minus current liabilities.	–
Non-Current Liabilities	NCLI	Total non-current liabilities.	–
Operating Profit/Loss c.	OPPL_calc	Operating profit (loss).	$\mathtt{OPPL\_calc} = \mathtt{OPRE} - \mathtt{COST\_shock}$
Profit Margin c.	PRMA_calc	Profit before tax / operating revenue .	$\mathtt{PRMA\_calc} = \frac{\mathtt{OPPL\_calc}}{\mathtt{OPRE\_adj}} \times 100$
Profit/Loss After Tax c.	PLAT_calc	Profit after tax (incl. minority interests).	$\mathtt{PLAT\_calc} = \mathtt{PLBT\_calc} - \mathtt{TAXA\_calc}$
Profit/Loss Before Tax c.	PLBT_calc	Operating profit + financial profit.	$\mathtt{PLBT\_calc} = \mathtt{OPPL\_calc} + \mathtt{FIPL}$
Return Shareholders' Funds c.	RSHF_calc	Profit before tax / shareholders' funds .	$\mathtt{RSHF\_calc} = \frac{\mathtt{OPPL\_calc}}{\mathtt{SHFD\_adj}} \times 100$
Return Total Assets c.	RTAS_calc	Profit before tax / total assets .	$\mathtt{RTAS\_calc} = \frac{\mathtt{OPPL\_calc}}{\mathtt{TOAS\_adj}} \times 100$

Continued on next page

Table AIII. Definitions of Variables Used in the Default Prediction Model (continued)

Variable name	Code	Description	Formula (if calculated)
Shareholders' Funds c.	SHFD_calc	Equity (capital + other shareholders' funds).	$\texttt{SHFD\_calc} = \texttt{CAPI} + \texttt{OSFD} + \Delta\texttt{PLAT}$
Solvency Ratio c.	SOLR_calc	Shareholders' funds / total assets .	$\texttt{SOLR\_calc} = \frac{\texttt{SHFD\_adj}}{\texttt{TOAS\_adj}} \times 100$
Taxation c.	TAXA_calc	Total tax charge for the year.	$\texttt{TAXA\_calc} = \begin{cases} 0, & \text{if } \texttt{PLBT\_calc} < 0 \\ \texttt{PLBT\_calc} \times \texttt{TAX\_RATE}, & \text{otherwise} \end{cases}$
Total Assets c.	TOAS_calc	Sum of non-current and current assets.	$\texttt{TOAS\_calc} = \texttt{FIAS} + \texttt{CUAS} + \Delta\texttt{PLAT}$
Total Shareholders' Funds c.	TSHF_calc	Total shareholders' funds.	$\texttt{TSHF\_calc} = \texttt{SHFD\_calc} + \texttt{NCLI} + \texttt{CULI}$
Working Capital	WKCA	Current assets minus current liabilities.	–
Year	year	Calendar year.	–

Table AIV. Input Variables for the Emission Prediction Model

Variable name	Code	Description
Cash and Cash Equivalents	CASH	Cash and cash equivalents.
Cash Flow	CF	Operating cash flow.
Category of Company	CATEGORY_OF_COMPANY	Company size class (SM, L, XL).
Creditor Days	CRPE	Average time to pay trade creditors (days).
Creditors	CRED	Trade and other payables.
Current Assets	CUAS	Total current assets.
Current Liabilities	CULI	Total current liabilities.
Current Ratio	CURR	Current assets / current liabilities .
Debtor Days	COLL	Average time to collect trade receivables (days).
Debtors	DEBT	Trade receivables.
Depreciation	DEPR	Depreciation and amortization expense.
Earnings Before Tax and Amortization	EBTA	Earnings before tax and amortization.
EBITDA Margin	EBMA	EBITDA as pct. of revenue.
Earnings to Turnover	ETMA	Earnings relative to turnover .
Employees (Number of)	EMPL	Headcount.
Extraordinary Items	EXTR	Income/expenses classified as extraordinary.
Finance Expenses	FIEX	Interest and other finance expenses.
Finance Profit/Loss	FIPL	Financial profit (loss).
Finance Revenue	FIRE	Interest and other finance income.
Fixed Assets	FIAS	Net fixed assets (after accumulated depreciation).
Gearing Ratio	GEAR	Long-term debt / shareholders' funds .
Gross Margin	GRMA	Gross profit as pct. of revenue.
Gross Profit	GROS	Revenue minus cost of sales.
Interest Cover	IC	Operating profit / finance expenses.
Interest Expenses	INTE	Interest expenses.
Intangible Fixed Assets	IFAS	Intangible assets.

Continued on next page

Table AIV. Input Variables for the Emission Prediction Model (continued)

Variable name	Code	Description
Legal Form	SLEGALF	Legal form / company type.
Listed on Stock Exchange	LISTED	Listing status (listed / not listed).
Liquid Ratio	LIQR	(Current assets – inventories) / current liabilities .
Loans	LOAN	Loans and borrowings.
Material Expenses	MATE	Raw materials and consumables used.
NACE Code (2-digit)	NACEPCOD2_2digit	Industry classification (NACE Rev. 2, two-digit).
Net Asset Turnover	NAT	Turnover relative to net assets.
Net Current Assets	NCAS	Current assets minus current liabilities.
Non-Current Liabilities	NCLI	Total non-current liabilities.
Operating Profit/Loss	OPPL	Operating profit (loss).
Operating Revenue	OPRE	Operating revenue.
Other Current Assets	OCAS	Other current assets.
Other Current Liabilities	OCLI	Other current liabilities.
Other Fixed Assets	OFAS	Other fixed assets.
Other Operating Expenses	OOPE	Other operating expenses.
Profit Margin	PRMA	Profit before tax / operating revenue .
Profit/Loss After Tax	PLAT	Profit after tax (including minority interests).
Profit/Loss Before Tax	PLBT	Profit before tax.
Profit/Loss for the Financial Year	PL	Profit for the year.
Provisions	PROV	Provisions.
Property, Plant and Equipment (PPE)	PPE	Property, plant and equipment.
Return on Assets	ROA	Net income / total assets .
Return on Capital Employed (pt)	ROCE	Profit before tax + interest, over capital employed.
Return on Capital Employed (at)	RCEM	Net income + interest, over capital employed .
Return on Equity	ROE	Net income / shareholders' equity .
Return on Shareholders' Funds	RSHF	Profit before tax / shareholders' funds .

Continued on next page

Table AIV. Input Variables for the Emission Prediction Model (continued)

Variable name	Code	Description
Return on Total Assets	RTAS	Profit before tax / total assets .
Share Capital	SCT	Issued share capital.
Shareholders' Funds	SHFD	Shareholders' equity (capital plus other shareholders' funds).
Shareholders' Funds to Long-Term Liab.	SOLL	Shareholders' funds / long-term liabilities .
Shareholders' Funds to PPE	SFPE	Shareholders' funds / PPE.
Shareholders' Funds to Total Assets	SHLQ	Shareholders' funds / total assets .
Solvency Ratio	SOLR	Shareholders' funds / total assets .
Staff Expenses	STAF	Staff costs (wages, salaries, social charges).
Stock Turnover	STOT	Inventory turnover.
Stocks	STOK	Inventories.
Tangible Fixed Assets	TFAS	Tangible fixed assets.
Total Assets	TOAS	Total assets (non-current + current).
Total Assets to PPE	TAPE	Total assets / PPE.
Total PPE	TPE	Total property, plant and equipment.
Total Shareholders' Funds and Liab.	TSHF	Total shareholders' funds and liabilities.
Turnover	TURN	Net sales.
Working Capital	WKCA	Current assets minus current liabilities.
Working Capital to PPE	WCPE	Working capital / PPE.
Year	year	Calendar year.
Year of Incorporation	YEARINC	Year of incorporation.
CO2 Equiv. Emissions Scope 1 (target)	`target`	Direct greenhouse-gas emissions, Scope 1 (CO_2e).
Country (ISO Code)	CTRYISO	Country of incorporation (ISO code).

B. Additional Figures

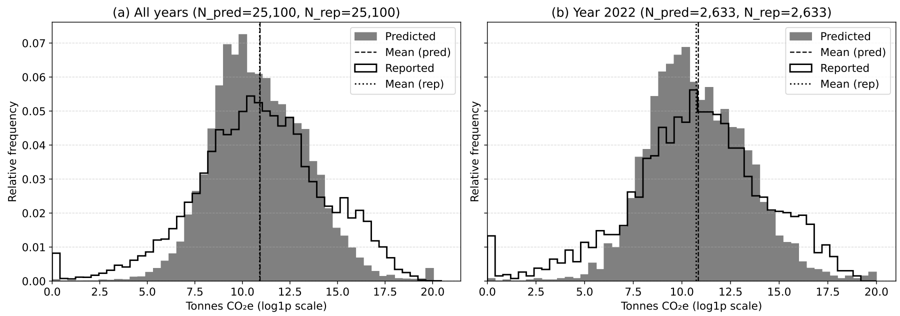

Notes: This figure presents the distribution of conditional mean predictions of firm-level Scope 1 GHG emissions and the distribution of actual reported emissions for a sub-sample of firms for which reported data are available. For the prediction, a Ridge-based model is used. Panel (a) presents the distribution for all years in the sample and panel (b) for the year 2022, which is used in the stress testing exercise.

Figure B1. Distribution of Predicted vs. Reported Scope 1 GHG Emissions for Reporting Firms (Ridge)

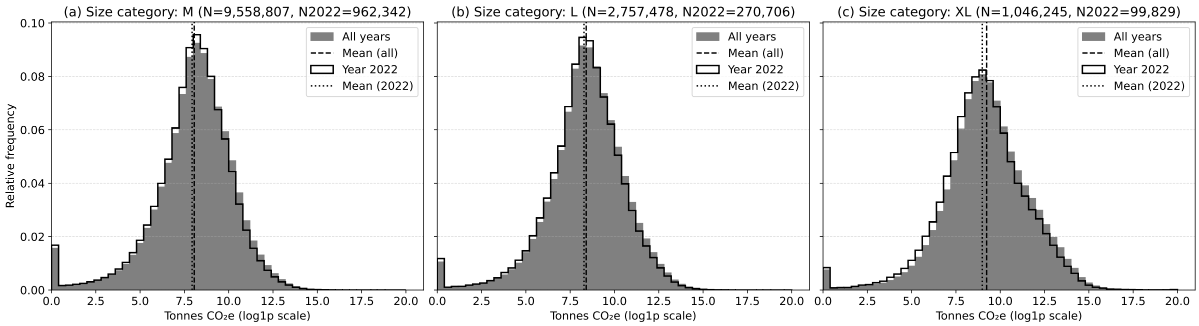

Notes: This figure shows the distribution of estimated Scope 1 GHG emissions for different firm sizes: panel (a) for M-sized firms, panel (b) for L-sized firms, and panel (c) for XL-sized firms. Firm size is defined as follows: *Very large (XL)*: at least one of revenue $\geq$ 100 m EUR, assets $\geq$ 200 m EUR, employees $\geq 1,000$, or listed; *Large (L)*: at least one of revenue $\geq$ 10 m EUR, assets $\geq$ 20 m EUR, or employees $\geq$ 150, not XL; *Medium (M)*: at least one of revenue $\geq$ 1 m EUR, assets $\geq$ 2 m EUR, or employees $\geq$ 15, not L or XL. Each panel shows the distribution for all years as well as for the year 2022, which is used in the stress testing exercise. The underlying prediction model is Ridge.

Figure B2. Distribution of Predicted Scope 1 GHG Emissions for Reporting Firms across Firm Sizes (Ridge)

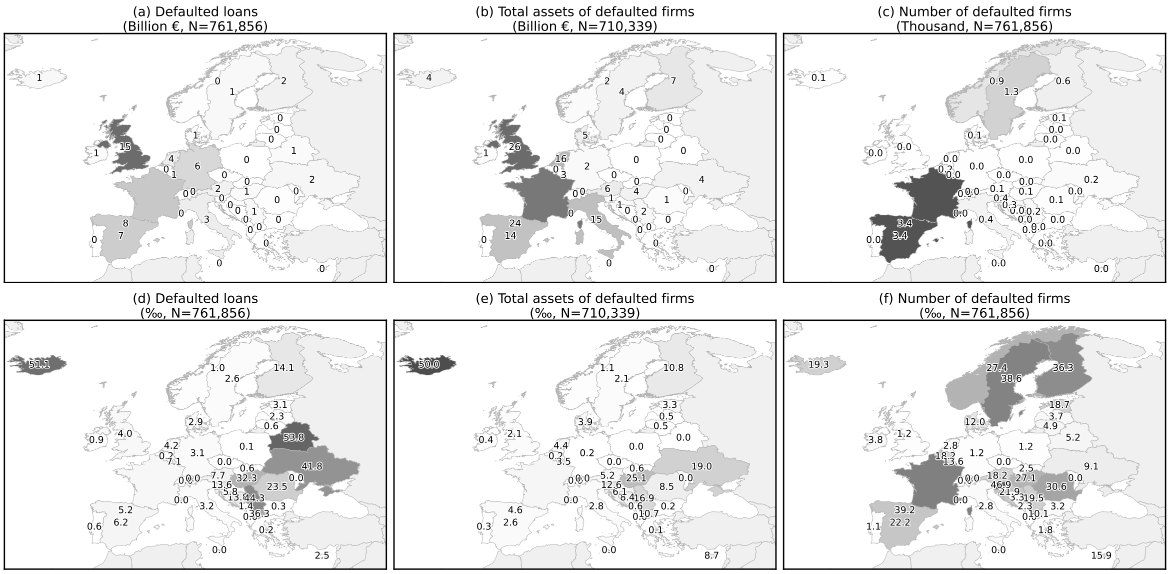

Notes: This figure reports the aggregate effects of a carbon price shock on firm defaults by country. Panel (a) reports aggregate defaulted loan volume, Panel (b) reports aggregate total assets of defaulted firms, Panel (c) reports the number of defaulted firms, Panel (d) reports the relative defaulted loan volume, Panel (e) reports the relative total assets of defaulted firms, and Panel (f) reports the relative number of defaulted firms after a carbon price shock ($\Delta CO_2\ price = 100\ EUR/t$) and in the baseline scenario ($\Delta CO_2\ price = 0\ EUR/t$) across countries.

Figure B3. Impact of a Carbon Price Shock on Firm Defaults by Country

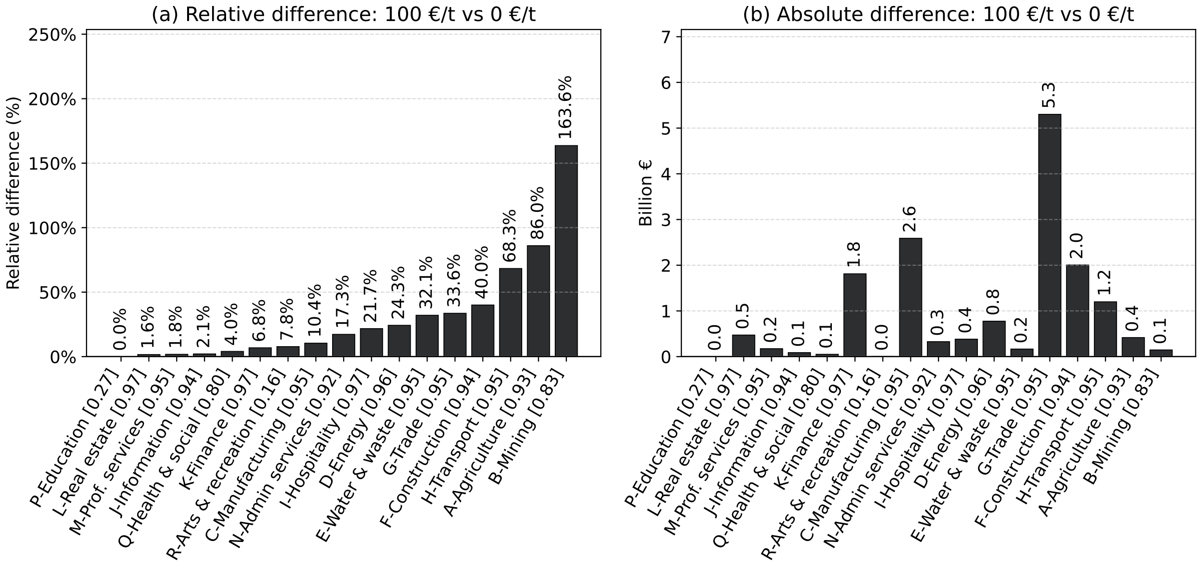

Notes: The figure reports the difference between total assets of defaulted firms after a carbon price shock (ΔCO_2 $price$ = 100 EUR/t) and in the baseline scenario (ΔCO_2 $price$ = 0 EUR/t) across sectors, as defined by NACE Rev. 2.1 level 1. Panel (a) reports the relative difference of total assets of defaulted firms. Panel (b) reports the absolute difference of total assets of defaulted firms. Sectors are sorted by relative differences. Due to missing data on total assets, the coverage (reported in square brackets) of firm observations as reported in Figure 11, varies by NACE section.

Figure B4. Defaulted Total Assets by Sector

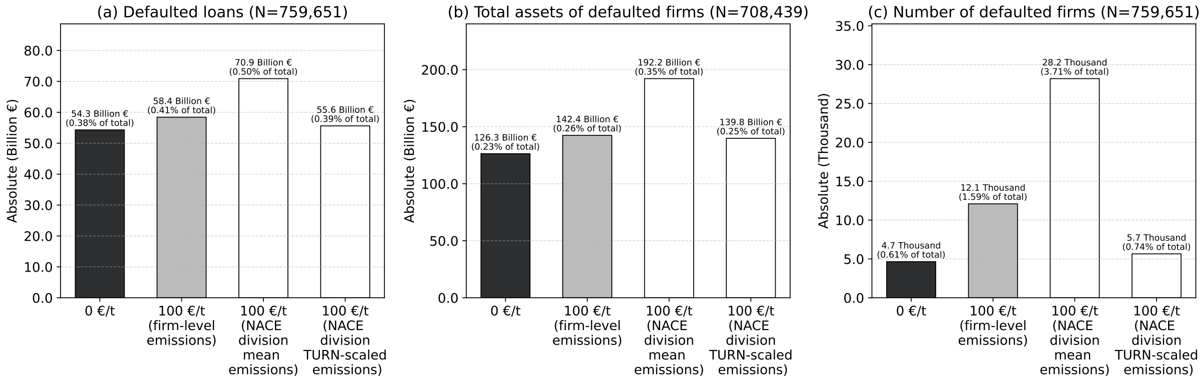

Notes: The figure reports the aggregate effects of a EUR 100 carbon price shock on firm defaults. Panel (a) reports the aggregate defaulted loan volume, Panel (b) reports aggregate total assets of defaulted firms, and Panel (c) reports the number of defaulted firms for a carbon price shock ($\Delta CO_2\ price = 100\ EUR/t$) and the baseline scenario ($\Delta CO_2\ price = 0\ EUR/t$). In all panels, the effect of the carbon price shock is computed based in predicted firm-level emissions, sector-average imputed emissions, and revenue-weighted sector-average imputed emissions, respectively.

Figure B5. Defaults with Predicted Firm Emissions versus Sector Averages