BANCA D'ITALIA
EUROSISTEMA

Questioni di Economia e Finanza

(Occasional Papers)

Forecasting corporate default probabilities: a local logit approach for scenario analysis

by Giuseppe Cascarino, Federica Ciocchetta, Stefano Pietrosanti and Ivan Quaglia

February 2025

Number 909

BANCA D'ITALIA
EUROSISTEMA

Questioni di Economia e Finanza

(Occasional Papers)

Forecasting corporate default probabilities: a local logit approach for scenario analysis

by Giuseppe Cascarino, Federica Ciocchetta, Stefano Pietrosanti and Ivan Quaglia

Number 909 – February 2025

The series Occasional Papers *presents studies and documents on issues pertaining to the institutional tasks of the Bank of Italy and the Eurosystem. The* Occasional Papers *appear alongside the* Working Papers *series which are specifically aimed at providing original contributions to economic research.*

The Occasional Papers *include studies conducted within the Bank of Italy, sometimes in cooperation with the Eurosystem or other institutions. The views expressed in the studies are those of the authors and do not involve the responsibility of the institutions to which they belong.*

The series is available online at www.bancaditalia.it .

ISSN 1972-6643 (online)

Designed by the Printing and Publishing Division of the Bank of Italy

FORECASTING CORPORATE DEFAULT PROBABILITIES: A LOCAL LOGIT APPROACH FOR SCENARIO ANALYSIS

by Giuseppe Cascarino*, Federica Ciocchetta*, Stefano Pietrosanti* and Ivan Quaglia*

Abstract

We propose a new approach for predicting corporate default probabilities and for conducting scenario analyses by combining firm-level and macro time series data. We apply a local projection approach to a simple logit framework and bridge the gap between micro data on firms, for which no scenario is available, and macroeconomic variables, for which the forecaster instead has a scenario. We apply this model to an out-of-sample exercise, estimating it with data through the end of 2017 and forecasting corporate defaults over the following three years. We compute two sets of projections, the first based on the realized values of the macroeconomic time series (baseline), and the second conditional on a scenario that simulates a worsening in the macroeconomic environment comparable to the one observed during the European sovereign debt crisis (adverse). The baseline forecast closely matches the actual corporate debt default rate; under the adverse scenario, the default rate is similar to the one actually recorded in Italy during the sovereign debt crisis. We also run two exercises that make use of the granular forecasts of the corporate default probabilities. First, we assess which sectors are more vulnerable under each of the previous two scenarios (baseline and adverse). Second, we assume that the economy shifts from the baseline to the adverse scenario and construct transition matrices across different risk classes, showing which sectors are more exposed to the shift.

JEL Classification: C25, C53, C54, G33.

Keywords: scenario analysis, logit model, credit risk.

DOI: 10.32057/0.QEF.2025.909

* Bank of Italy, Financial Stability Directorate.

1. Introduction[1]

Firms defaulting on their bank loans pose a critical risk to the stability of the financial system and the economy at large. Defaults correlate with the accumulation of non-performing loans on banks' balance sheets, which impairs banks' capacity to lend per se (Accornero et al., 2017) and may eventually drive banking crises, with their dire consequences (Reinhart and Rogoff, 2009; Schularyk and Taylor, 2012). Firms close to default may increase their borrowing, displacing healthier firms and increasing impairments in the future (Ivashina and Scharfenstein, 2010; Giesecke et al., 2014). Thus, understanding the drivers of corporate defaults and predicting their evolution under different economic conditions are important for monitoring financial stability risks.

Most models employed by central banks and regulatory authorities focus on aggregate data and forecast the average corporate default rate. These models can perform conditional forecasting given a macro scenario (e.g. Bonaccorsi di Patti and Cascarino, 2020; Gross and Población, 2017). On the contrary, the corporate finance literature generally studies the cross-section of individual corporate defaults using firm-level factors (for a survey of classic empirical works studying defaults at the firm-level, see Altman and Hotchkiss, 2006).

We propose a modelling approach to obtain firm-level projections of loan default rates combining microeconomic information (balance sheet and Bank of Italy's Central Credit Register data) with macroeconomic variables (real and nominal).[2] The approach allows us to build a flexible toolkit for scenario analysis. In particular, this model can be used (i) to assess the evolution of default risk for groups of firms (industrial sectors, geography, size-classes, types of lenders) especially exposed to unforeseen shocks; (ii) to simulate the impact on credit risk of firm-targeted policies, while accounting for various heterogeneity dimensions; and (iii) to perform stress testing.

A key challenge that we must address to construct a mixed micro-macro model that performs forecasting over a multi-year horizon is how to combine data updated with different frequencies and available after different lags (shorter for the macroeconomic variables; longer for balance sheet information). Most importantly, projections are typically available only for macroeconomic variables. As the forecast horizon moves further in the future, the microeconomic data age, while the macroeconomic data stay current because of scenario projections' updates. Thus, we cannot simply fit forward the same model.

We adapt the logic of local projections (Jordà, 2005)[3] to address this challenge; thus, we estimate a different specification for each forecasting horizon. We first define the one-year-ahead default equation and then, to perform the h-years-ahead forecast, we estimate an equation that exploits the latest firm-level information up to the last year available in the sample, minus h. Thus, we forecast defaults exploiting the data available at each forecasting horizon.

Each h-years-ahead forecast equation is a logit hazard model in the style of Shumway (2001); thus, it predicts the probability of default of each firm. Starting from a set of microeconomic and macroeconomic variables, we identify the best one-year ahead equation using a selection algorithm that optimizes forecast performance. First, we identify the specifications with the highest performance in discriminating between future defaulters and survivors; second, among the specifications selected

[1] We would like to thank Alessio Anzuini, Emilia Bonaccorsi di Patti, Alessio De Vincenzo for their helpful comments and suggestions.

[2] Other recent works explored firm-level and macroeconomic default predictors jointly, e.g. Jacobson, Lindé, and Roszbach (2013); Filipe, Grammatikos, and Michala (2016). Nevertheless, these works do not tackle scenario analysis.

[3] Jordà (2005) first introduced local projections to recover impulse responses from a sequence of distinct regressions at different horizons, instead of projecting the same multivariate dynamic model over progressively longer horizons (as in VARs). The local projections method is highly flexible and has been applied to a variety of topics, such as the effects of geopolitical risk (e.g. Caldara and Iacoviello, 2022), of credit booms (e.g. Jordà, Schularick, and Taylor, 2016), and how these and other kinds of shocks interact with firms' heterogeneity (e.g. Crouzet and Mehrotra, 2020; Fabiani, Falasconi, and Heineken, 2022; Ottonello and Weinberry, 2020).

in the first step, we choose the best at fitting the average default dynamics. We then estimate the two-year-ahead and three-year-ahead equations by increasing by one and two years the distance between the micro and macro covariates. Thus, we keep the resulting specification fixed for the two and three years ahead horizons. In conclusion, our model is the three independently estimated equations.

The Granular and Macroeconomic Probability of Default Model we present in this draft (GRANMA-PD model henceforth) performs very well in terms of in-sample fit and can correctly point out defaulters out-of-sample at each horizon. The predicted default probabilities of failing firms are systematically higher than those of surviving firms. The Area Under the Receiver Operating Characteristic curve (AUROC) amounts to about 80 percent for all equations, both in-sample and out-of-sample, which confirms a good discrimination capacity, matching the performance of the machine learning model by Cascarino, Moscatelli, and Parlapiano (2022).[4]

Studying a set of alternative specifications that omit groups of predictors, we find that the macroeconomic variables' main contribution is fit. Without macroeconomic variables, the model could not match the average default rate dynamics, i.e. the overall level of risk. In contrast, firm-level data increase the model's capacity to single out future defaulters. These findings are consistent with the results in Jacobson, Lindé, and Roszbach (2013).

We show two applications of the GRANMA-PD model. In the first, we collapse our firm-level default forecasts to obtain an average default rate and perform an out-of-sample exercise. We estimate the model with data up to 2017 Q4 and forecast defaults over the years 2018, 2019 and 2020 under two different scenarios. The baseline scenario takes as inputs the realized values of the macroeconomic time series, while the adverse scenario simulates a worsening in the macroeconomic conjuncture comparable to the one recorded in Italy during the European Sovereign Debt crisis. The resulting average of the baseline forecasts matches the realized corporate debt default rate closely, while the average of the adverse forecasts implies a counterfactual 30 percent increase in the default rate, close to the one experienced from 2011 to 2012.

While our forecast period ends with the onset of the COVID-19 pandemic, our baseline exercise does not foresee jumps in the average default probability. The lack of foreseen structural breaks should not come as a surprise for two reasons. First, the macroeconomic time series informing our forecasts end at Q4 2019. Thus, the macroeconomic effects of the pandemic are excluded from our inputs, making our model "COVID-ignorant". Second, the macroeconomic conjuncture did not react immediately to the pandemic crisis, while the moratorium on corporate debt and other aid limited the immediate impact of business closures.[5] In conclusion, we can interpret our baseline forecast for year 2020 as projecting the average default rate in the absence of the pandemic; the good fit of the model even in that year provides evidence that support measures kept corporate debt default on its initial path at least up until the end of year 2020.[6]

[4] We also test that the high AUROC is not the result of forecasting on an unbalanced sample. We focus on firms' defaults and these are always a minority in the data. Therefore, a bogus model always forecasting survival could appear to be "good" from a naïve reading of AUROC-like indicators. To address this possible pitfall, we also compute the level of Specificity (the capacity to correctly foresee survivors) and Sensitivity (the capacity to correctly foresee defaulters) of our forecasts using as a default threshold on the fitted probabilities equal to the average default rate in the whole of the sample. We find a Sensitivity of 78 percent and a Specificity of 67 percent. These numbers point to the fact that our good performance comes exactly from the correct forecasting of defaulters, thus mitigating the unbalanced sample concern.

[5] First, the Decree Law 18/2020 ('Cura Italia') introduced a debt moratorium targeting small-to-medium-sized enterprises (SMEs) that had no non-performing loans when the measure was introduced (see the 'Public intervention in lending to firms' Box within the Bank of Italy Financial stability Report, n.1 2020). Furthermore, the Decree Law 23/2020 ('Liquidità') introduced the opportunity for firms to access new state-backed loans via the Central Guarantee Fund or SACE.

[6] Multiple works studied the role of the wide array of economic support measures implemented by the Italian government to limit corporate bankruptcy liquidations during Covid pandemic, reaching similar conclusions. See, for instance, Giacomelli *et al.* (2021) and Orlando and Rodano (2022). Emergency support measures enabled many companies to cover, at least partially, their liquidity needs and potential equity deficits (see also De Socio *et al.*, 2020, Orlando and Rodano,

The second application uses the firm-level granularity of the estimates to map the distribution of default probability forecasts for all firms in each sector. We propose two examples of how we can use such granular information. First, we look at how the distribution of firms' default probability forecasts varies across industries fixing the macroeconomic scenario. Such an exercise helps tracking divergence in the sectoral evolution and concentration of risk under a determinate conjuncture. Then, we compute the share of firms that transition between four risk classes, based on the probability of default value, under the adverse scenario. This last exercise provides us more detailed insight into the drivers of fragility, pointing out whether a worsening of the situation of the core of healthy firms drives the increase in the average forecast or whether this is due to the tail of firms whose probability of default was already high under the baseline scenario.

The remainder of the paper is as follows. Section 2 discusses the literature on firms default forecasting and its relation to our approach. Section 3 describes the dataset used in our work. Section 4 presents the methodology and the results of our average default forecasting exercise. Section 5 illustrates the two applications exploiting the granularity of our forecasts to track the evolution of firm default risk within and across scenarios. Finally, Section 6 highlights the main conclusions of the work.

2. The literature on default risk quantification

There are two established approaches to quantify firm-level default risk: first, risk scoring using structural models of the stochastic movement in firms' value of assets; second, statistical models of default, either based on discriminant analysis or on logit regressions fitted to historical firm-level data. Recently, however, the field has been undergoing a wave of innovations. Applications of machine learning and artificial intelligence promise to help optimize the variable selection step, while improving forecasting performance. We provide here a brief review of these strands, to better frame our contribution to the subject matter.

The structural approach descends from Merton (1974). Merton's paper builds a forward looking Distance to Default (DD) model. The underlying theory is that firms default if their assets' value falls below their liabilities' value. The model evaluates assets and liabilities based on market price data, derives the difference, and finally fits the relationship between this difference and data on defaults (see Altman and Hotchkiss, 2006, for a synthetic description). This framework is extensively used in valuation (see the recent review in Bakshi, Gao, and Zong, 2022), but it requires information on the market value of firms' assets and liabilities, which is usually not available for small firms. This limitation is especially relevant to us, as 64 percent of Italian value added and 75 percent of employment is due to SMEs, making Merton's model unviable.[7]

Lack of market value of equity implies that applying Altman (1968)-like discriminant analysis is not possible for us either, as the market value of equity is part of the inputs in Altman's baseline approach. We thus build on logit-based approaches, similarly to Shumway (2001), estimating a hazard model of firm default. Hazard models explicitly deal with the passing of time in discrete steps, allowing us to tackle the relationship between default and the evolution of firm and macroeconomic conditions. We modify this framework by including industry and macroeconomic variables, thus our work is also related to contributions such as Chava and Jarrow (2004), showing the relevance of industry effects, and Jacobson, Lindé, and Roszbach (2013), Filipe, Grammatikos, and Michala (2016) and Jensen, Lando, and Medhat (2017), focusing on the importance of the economic conjuncture.

In as much as we include an automated variable selection step in the design of our default model, we are related to the longstanding line of works applying machine learning to corporate default forecasting (e.g. Altman, Marco, and Varetto, 1994). Similar approaches are gaining new momentum

---

2020, and the box 'The effects of the pandemic on the balance sheets and riskiness of firms in the various economic sectors', in Bank of Italy, Financial stability report No. 1 – 2021).

[7] The figures quoted come from the European Commission Italian SMEs' fact sheet 2024, collected within the Small Business Act and available at https://ec.europa.eu/docsroom/documents/60571

thanks to the greater availability of computational power and to novel software advancements. Recent applications show that these techniques can achieve improvements over basic logit models (e.g. Bachman and Zhao, 2017; Barboza, Kimura, and Altman, 2017; Moscatelli et al., 2020). Nevertheless, there is an accuracy *vs.* interpretability trade-off, as such models often lack an easily interpretable functional form (Cascarino, Moscatelli, and Parlapiano, 2022).

Finally, the GRANMA-PD model complements other tools already employed by the Bank of Italy's Financial Stability Directorate to assess risks to financial stability arising from nonfinancial firms. First, the model's firm-level forecasts can be aggregated at will, in contrast with time-series models, such as Bonaccorsi di Patti and Cascarino (2020), whose level of aggregation is pre-set. Second, the model also complements the In-house Credit Assessment System (ICAS, Giovannelli et al., 2020). The ICAS model forecasts the probability of default at the firm level too. However, it does so only at the one-year ahead horizon and does not include macroeconomic variables directly.[8] As the GRANMA-PD model forecasts firms' default at multiple horizons and includes such macroeconomic variables, it can be used for stress testing and scenario analysis. Finally, the GRANMA-PD model can be integrated with the microsimulation model by De Socio and Michelangeli (2017). De Socio and Michelangeli (2017) simulates firms' balance sheets over a two-year horizon to assess financial vulnerability under given scenarios, rather than default probabilities.

3. Data and variables

To build our forecasting model, we draw on the following sources:

1. Credit information from the Central Credit Register (CCR), updated monthly with a time lag of two months.
2. Corporate balance sheet data from the Company Accounts Data system (provided by Cerved) available on an annual basis, with a lag of about one year.
3. Macroeconomic quarterly data and projections updated by the Bank of Italy's scenario exercise for the Italian economy for the next three years.[9]

Our target variable is the probability that a firm will be in default status over a certain time horizon, conditional on being initially healthy. For each quarter q in our dataset, and for each non-defaulted firm in q, we define as the dependent variable a dummy indicator (the new default dummy) equal to 1 if the firm is in default in year y_0+h starting from quarter q. We consider h equal to one, two or three years. For the time horizon of 2 years, we also impose that the firm is not in default at time y_0+1 and in the case of 3 years, that the firm is not in default at both time y_0+1 and y_0+2.

We employ the same default definition as for the system-wide non-performing status of a borrower, reported by the Bank of Italy in its Central Credit Register. A firm is classified as being in default at a given date of reference if the ratio of non-performing to total credit drawn from the banking system is greater than five per cent for at least one month. We can aggregate this measure the average default rate, i.e. the fraction between new non-performing borrowers over a one-year period and performing borrowers at the beginning of that period, calculated quarterly. Such average ratio is an aggregate measure of credit risk, which will constitute our first validation target.

We describe the behavior of the average default rate over the last 16 years in Figure 1. There, we can observe a twin-peaks evolution, with highest levels hit in 2009, after the Great Financial Crisis, and

[8] This last model's main objective is to assign a risk rating to Italian firms. This rating enables banks to use loans granted to these firms as collateral for the Eurosystem.

[9] The macroeconomic projections for Italy are released in April, June, October and December on the Bank of Italy's website. On exceptional basis, an update of the macroeconomic projections for Italy can be released to take account of new data and of possible changes in the external environment. For more details see: https://www.bancaditalia.it/pubblicazioni/proiezioni-macroeconomiche/index.html.

2014, in the aftermath of the European sovereign debt crisis with the associated slowdown of the Italian economy. After that, the aggregate default risk started decreasing, with default rate just above 1% at the end of the observed period. A good risk model, one that we can use to project this series forward and perform scenario analysis, must be able to replicate the default behavior over time both within and out of sample.

Figure 1: The average default rate over the last 16 years

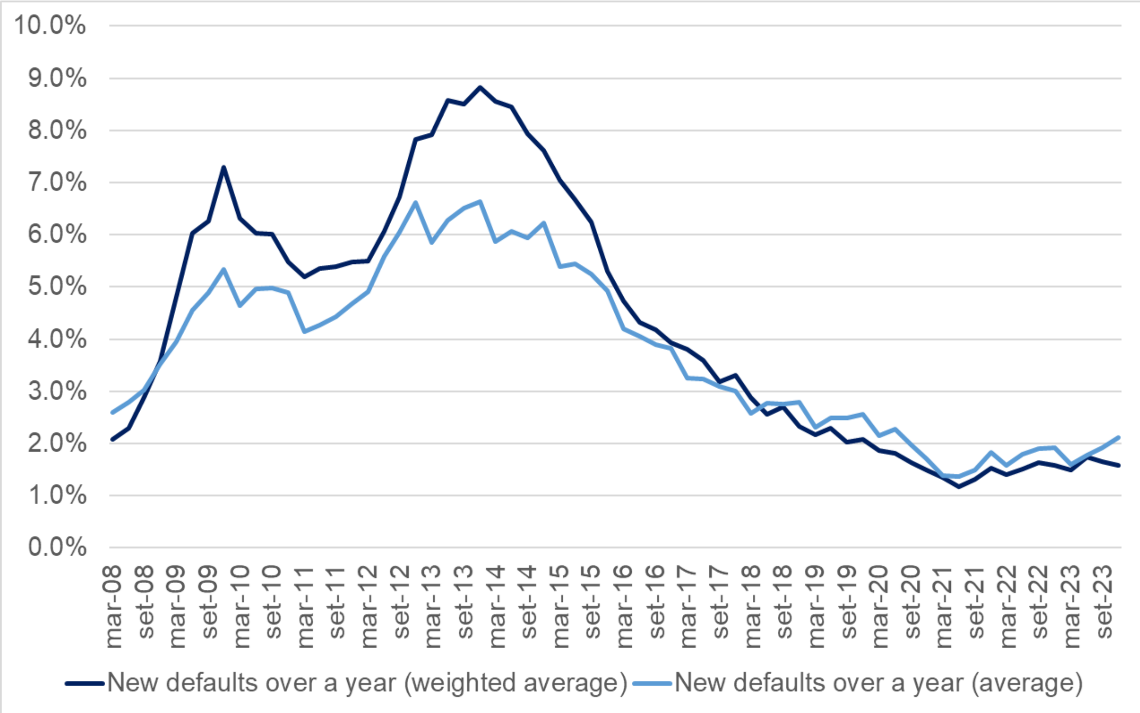

Sources: Central Credit Register. *Note*: The annual average of new defaults is the proportion of firms in our sample that will be in default in a one-year period over the number of firms not in default at the beginning of that period, calculated quarterly. For example, mar-08's data point covers the default rate between 2007Q1 and 2008Q1. The weighted average of new defaults is the same indicator, but weighted by the (drawn) loan amount.

As regards the possible regressors in the model, we rely on balance sheet and credit indicators for Italian non-financial firms (micro variables) and macroeconomic data. We base a first selection of the variables on the literature on corporate default and bankruptcy risk (especially internal works using the same sources, e.g. Bonaccorsi di Patti and Cascarino, 2020; Moscatelli et al., 2020). Micro variables measure profitability, financing structure, debt sustainability, credit drawn and granted amounts for non-financial corporations, all at the firm level.

We then select as macro variables five components of the macroeconomic conjuncture whose relationship with insolvency we deem of primary importance (see Table A.1):

1. The growth rate of real GDP, as a proxy of demand movements that may push or depress firms' profitability and, as a consequence, firms' debt repayment capacity.
2. The average nominal interest rate (cost of credit) on new loans paid by non-financial corporates, which directly affects corporates' debt repayment ability. We consider this important, first, because the weight of debt service rises with loan rates, directly increasing firms' vulnerability. Second, because loan rates are tightly related to bank funding conditions. Rates' systemic increase often signals negative credit supply shocks, which harm firms in need to refinance their loans (Bonaccorsi di Patti and Cascarino, 2020).
3. Inflation, as measured by the growth rate of the national final consumption deflator. As we measure the cost of credit in nominal term, it is important to also control for changes in the value of money. Moreover, as observed in Jacobson, Lindé, and Roszbach (2013), the overall effect of inflation on credit risk is not obvious. Indeed, increases in nominal profitability may increase corporate debt repayment capacity if fixed rate debt is a quantitatively important

source of funding. Nevertheless, rise in inputs' cost and price uncertainty may harm sales and profits, limiting firms' debt repayment capacity. Moreover, demand and supply driven inflation episodes may have different effects. Indeed, prices rising for increasing demand imply also positive news for producers, while, if supply bottlenecks or cost push shocks drive inflation up, firms may expect only a decrease in margins.

4. The growth rate of house prices. Residential real estate is a primary source of collateral for European firms, with 30 percent of firms with less than 50 employees reporting pledges of entrepreneurs' personal assets, including the owners' home (Banerjee and Blickle, 2016; data from SAFE survey).
5. The firm gross operating margin (EBITDA). We use an index number summarizing the information across all firms in the Italian economy. This indicator measures the efficiency of the non-financial corporate sector.

The four macroeconomic variables are updated quarterly, and their projections one, two and three years ahead are computed at least four times a year with the Bank of Italy's quarterly econometric model.[10]

To obtain the dataset on which we estimate and project our models, we convert all data to a quarterly frequency. We repeat over each quarter data that is updated yearly, and disregard mid-quarter observations for data that is updated monthly. We focus on the quarterly frequency to align with the quarterly scenario's update frequency. The resulting dataset includes 840 thousand distinct firms and almost 300 thousand firms per quarter.

4. Risk assessment framework: A local logit approach

We face a fundamental challenge in building a multi-period forecast model for the probability of default at the firm level. Firms' balance sheet data are only available yearly with a one-year lag. Moreover, these data are static, with no forward projection. Macroeconomic data are instead dynamic, as the Bank of Italy updates its forecast quarterly. Consequently, we cannot just forward the one-year-ahead model to obtain the two- and three-year-ahead predictions. Indeed, as we forecast firms' default probability further in the future, the lag between the micro and macroeconomic information would widen, making our estimated coefficients incoherent with the underlying data. To address this problem, we estimate three different binary logit models to forecast firm defaults at the one, two, and three-year horizons, inspired by the local projection approach (Jordà, 2005). Each model is local to its forecast horizon, thus exploiting information as well as possible in its timeframe.[11]

We choose the logit model as our building block based on the result in Shumway (2001), showing that a multi-period logit model is equivalent to a hazard model with a logistic cumulate density, describing survival probability up to time t. Hence, it is a natural framework to predict the probability that a firm defaults in time $t+1$, conditional on having survived up to t.

[10] The macroeconomic projections for the euro area are released on the web site of the European Central Bank in March, June, September, and December of each year. In June and December these projections are produced on a set of common assumptions by experts from the Eurosystem national central banks, jointly with ECB staff members. Twice yearly (in June and December) the Bank of Italy releases a brief comment on the macroeconomic projections for Italy. A more detailed analysis of the outlook for the Italian economy, including any updates made necessary by changes in the external environment or by new data, is published in the January and July issues of the Bank of Italy's Economic Bulletin. Furthermore, updated scenario can be implemented if necessary for internal use.

[11] Such approach is conservative, for simulating firm-level balance sheet variables is a complex task and requires taking a stance on the drivers of firms' balance sheet evolution. The upside of our simple approach is not to make such further assumptions, keeping the moving parts at a minimum. The downside is a limitation in capturing possible feedback loops from macro to micro. To address such limitation in the future, we are also working on an extension of the model that integrates simulated balance sheet variables from De Socio and Michelangeli (2017) microsimulation model.

We believe this to be the best parametric option. A valid alternative would have been the implementation of a bottom-up non-parametric model, for example the machine-learning-based framework in Cascarino, Moscatelli, and Parlapiano (2022). We favored starting from a parametric perspective as it allows greater transparency and interpretability of the results. Our framework is however amenable for integration with interpretable machine-learning approaches, which we leave for further research. Nevertheless, we do include a bottom-up component, i.e. an automated variable selection step to simplify and discipline our choice of macroeconomic variables for the first model (one-year horizon, details in Section 4.2).

4.1 Empirical model

We estimate the following linear model in the log-odds of default:

$$\log\left(\frac{p_{q,y0+h}}{1-p_{q,y0+h}}\right) = \mathrm{B}_{\mathrm{h}}^{\mathrm{CE}} X_{y0}^{micro,CE} + \mathrm{B}_{\mathrm{h}}^{\mathrm{CR}} X_{q,y0}^{micro,CR} + \mathrm{B}_{\mathrm{h}}^{\mathrm{macro}} X_{q,y0+h}^{macro}, \qquad h = 1,2,3 \qquad (1)$$

where q is the quarter of the year y_0 at which the default information was last observed, and h is the forecast horizon in years. $p_{q,y0+h}$ is the probability that a firm that did not default up to quarter q of year y_0 defaults by the same quarter of year y_0+h (also PD, henceforth); $X_{y}{}^{micro}{}_{0}{}^{,CE}$ is the matrix of the latest available balance sheet information, and $X_{q}{}^{micro}{}_{,y0}{}^{,CR}$ contains the credit information available at quarter q of the year y_0; $X_{q}{}^{macro}{}_{,y0+h}$ is the matrix of macroeconomic variables (possibly lagged), that reaches in the future thanks to the Bank of Italy's macroeconomic projections. Finally, $\mathrm{B}_{\mathrm{h}}^{\mathrm{CE}}$, $\mathrm{B}_{\mathrm{h}}^{\mathrm{CR}}$, $\mathrm{B}_{\mathrm{h}}^{\mathrm{macro}}$ are the parameter vectors we must recover.

To estimate and validate the above three independent logit equations we perform three steps:

1) Variable selection. We consider several combinations of micro and macro variables (with different lags) and, among them, we select the subset of covariates which optimizes the prediction performance.

2) Estimation. For each forecast horizon h, we estimate a logit model where the dependent variable is the new default dummy described in Section 3, and the regressors are the optimally chosen ones. We estimate each h-model on its own dataset, which encodes the appropriate structure of lags to reflect the different control variables' update pattern. Which is, take the qth quarter of year y, the balance sheet and credit register information concerns the same quarter of year y-1 for the model forecasting the one-year-ahead PD, y-2 for the two-years-ahead model, and y-3 the three-years-ahead model. The macro variables are, instead, always fully updated due to the scenario. In conclusion, the $\hat{\mathrm{B}}_{\mathrm{h}}^{\mathrm{macro}}$ vector will reflect the impact of the macroeconomic variables on default probability, net of all the information that can be extracted from micro data last updated h years ago. We obtain 95 percent confidence intervals for our predicted probabilities by bootstrapping. We detail the bootstrapping procedure in Annex B.

3) Validation. We compute prediction performance of the three models along two dimensions, (i) the ability to predict average default rates (for all firms and for industrial sectors) and (ii) the predicting power for future defaults at the firm level.

4.2 Variable selection algorithm

We identify the optimal set of variables among the ones reported in Tables A.1 and A.3. We do so targeting the best possible performance in terms of discriminating power between defaulted and non-defaulted firms and fit of the average default rates aggregated from the model predictions. We implement a two-step algorithm, based on two measures of performance, the Area Under the Receiver Operating Characteristics Curve (AUROC) and the Root Mean Square Error (RMSE). The AUROC measures the probability of correctly forecasting firms' default or survival;[12] its optimization ensures that our model assigns high default probability to future defaulters and vice versa. The RMSE is, instead, a standard measure of fit.[13] Targeting it, we make sure that fitted default probabilities aggregate to an average default rate close to the one we observe on a per-firm basis.[14]

In the first step of the algorithm, we estimate several one-year-ahead probability of default models, for different combinations of micro and macro variables.[15] Then, we compute the AUROC in-sample for each model, using a training dataset including observations up to 2016 and keeping the models with an in-sample AUROC of 0.8 or greater.

In the second step, we apply an automatic procedure based on backward search to find, among the models detected in step 1, the specifications that have a low RMSE for the out-of-sample aggregate default rate's prediction. From this set of specifications, we select the reference model by expert judgment, to account for the fact that a fully automated procedure may incur in local minima, and may fail to include key scenario variables.

Both steps are important. Indeed, a model with a great performance in discriminating between future defaulters and survivors (identified by AUROC) may yet be far from the best in fitting the average default rate dynamics. In practice, the first step selects credit register and balance sheet controls, whereas the second step the set of macro variables and their lags.[16]

The resulting set of covariates includes seven firm-level variables plus five macro variables, each of which may appear in our equations with multiple lags. In detail, the **micro variables** are:

1) Capitalization, measured by the ratio of equity over total assets, which tracks firms' capacity to absorb shocks with own resources (and is negatively correlated with leverage).
2) Liquidity, measured as the ratio of current assets and short-term debt (current ratio), which tracks each firm's ability to pay off short-term debt obligations with cash and cash equivalents.
3) The ratio of interest expenses over cash flow, indicating the enterprise's ability to pay interest promptly.
4) Financial flexibility, i.e. the proportion of total drawn to total granted bank credit to the firm, which tracks how each firm can absorb shocks by increasing its borrowing.
5) Profitability, measured as the ratio between gross operating margin (Earnings Before Interest Taxes Depreciation and Amortization, EBITDA) and total assets. The higher this ratio, the higher the firm capacity to generate cash flow from its assets and, consequently, to repay debt.
6) Dummies for firm size based on the European Commission definitions of micro, small, average and large firms.
7) Dummies for the economic sector in which the firm operates, which account for individual industry exposure to shocks, often an important determinant of revenue volatility and fragility (see Chava and Jarrow, 2004).

[12] In greater detail, the AUROC is a summary statistic for the likelihood of predicting the default of a firm that will not default, and vice-versa, under a given model. It varies between a half, characterizing a model that has no capacity to distinguish between future defaulters and future survivors, and one, a model that is perfectly able to distinguish between future defaulters and future survivors.

[13] To provide accurate average scenario forecast, the average fit tracked by the RMSE is important in its own sake, even if it is not crucial for discriminating defaulters and non-defaulters. This attention is coherent with standard risk-modeling calibration practices, where analysts verify the models' capacity to replicate in-sample default rates (e.g. Giovannelli et al., 2020).

[14] Albeit we focus on firm-level default, our fit is not sensitive to weighting firm defaults by each firm loan amounts.

[15] We considered possible combinations of 4-6 variables, on the basis of the evidence from the literature and not including indicators with similar meaning.

[16] The algorithm adapts standard practices in model selection: See, for example, Bruce and Bruce (2017) and James et al. (2014). In particular, the separation between micro and macro variable selection steps is common in default modeling: See, for example, Filipe, Grammatikos, and Michala (2016).

We report the list of the micro variables included in the dataset in Table A.2 and their descriptive statistics in Table A.3.

The **macroeconomic variables** we include are:

1) The spot and six months lag of the average nominal interest rate on new corporate loans.
2) The spot and six months lag of inflation.
3) The six months and yearly lag of the growth rate of real GDP.
4) The two years lag of the growth rate of house prices.
5) The spot value of firm gross operating margin (EBITDA) in the Italian economy.

4.3 Model performance and estimated average effects

In Table 1, we show the values of in and out-of-sample AUROC. First, from the fact that the out-of-sample AUROC is also high, we can see that we are not overfitting, and we can robustly point out risky firms. Second, we use the AUROC statistic to compare two versions of the one-year-ahead logit equation and show the role that different variables play. The first version is our baseline; we call it the "Reference model" henceforth. The second is a micro-data-only model, which we call "Firmvariable model". From the AUROCs' similarity in Table 1, we can appreciate that balance sheet and credit information are enough to make accurate firm-level default predictions.[17]

[17] We check that this accuracy is not due to always foreseeing survival of firms in an unbalanced sample in which defaults are always a minority of observations. To do so, we check our model performance in terms of specificity, which is, the likelihood of foreseeing no-default for a firm that will actually survive, and – more importantly - sensitivity, which is, the likelihood of foreseeing default for a firm that will actually default. In order to compute these statistics, we need to actually turn the estimated probability of default into zero-one default predictions. We follow standard practice and assign to the default state firms with an estimated default probability greater than the sample average. Doing so, we obtain a specificity of 67 and a sensitivity of 78 percent for the one-year-ahead model, and similar performances for the other two logit equations. In conclusion, our high AUROC does not come from being correct only concerning survivals in a sample in which survival is the most likely outcome.

Table 1: Balance sheet and credit variables are sufficient for firm-level default forecasts

Model specification	AUROC in-sample	AUROC out-of-sample
Reference model	0.803	0.803
Firm-variable model	0.801	0.802

Note: All the figures reported concern one-year-ahead default models. We estimate two such models: The Reference model is the baseline, selected through the algorithm described in Section 4.2. The Firm-variables model includes only micro variables. We estimate each model using data from 2007Q1 to 2015Q1. We perform the out-of-sample AUROC calculation using data from 2015Q2 to 2016Q4.

Figure 2: Balance sheet and credit variables are not sufficient to match default rate dynamics

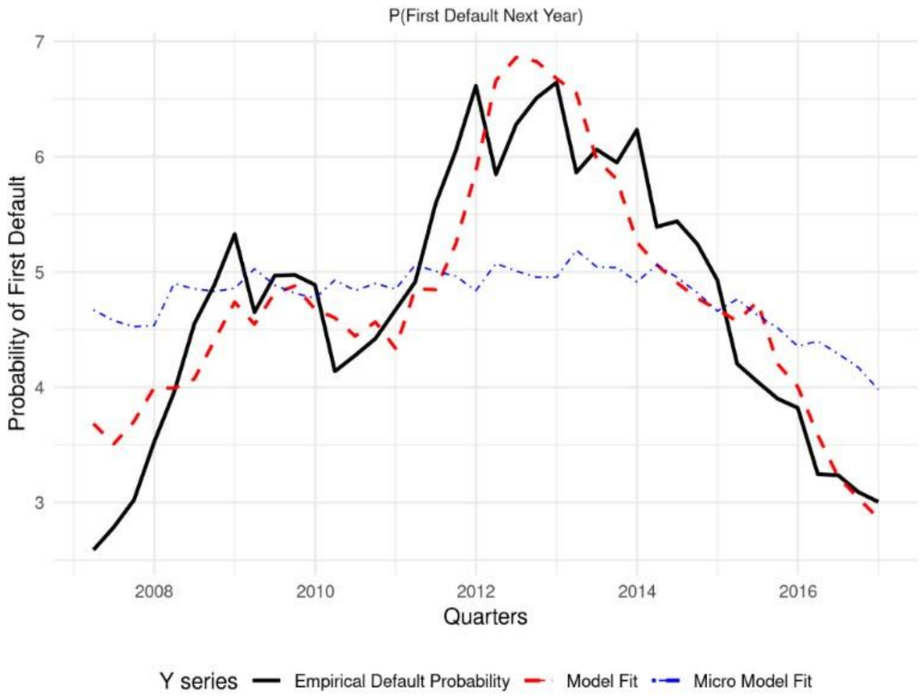

Source: Fit for the one-year-ahead default model. Authors' calculation based on Cerved and Credit register data. *Notes*: The Figure presents fits for Table 1's reference and firm-variables models, estimated on the whole sample up to 2016Q4. Time labels in the plot are moved forward by one year due to the default variable definition.

Nevertheless, in Figure 2 we show that we cannot match the average default rate dynamics without macro variables. We report the aggregate estimated defaults one-year-ahead, obtained by using our reference model and the variant that includes the same set of micro variables but no macro variables. We plot the fitted values against actual aggregate default rates.

While the two models have similar discriminatory power between defaulters and non-defaulters (in terms of AUROC, see Table 1), the model that relies only on micro variables does not yield an accurate estimate of the average rate of corporate defaults over time. The fit is instead very good for the model including both micro and macro variables. This result is consistent with the literature. For example, Jacobson, Lindé, and Roszbach (2013), studying logit models for Swedish firms' defaults, show that while an only micro model can still make a reasonably accurate ranking of firms' according to their riskiness, macroeconomic variables are key to explaining the average default rate dynamics.[18]

4.4 Estimation results

In Figure 3, we report the average marginal effects (AME) for each variable across the three models. AMEs are the standard statistic to summarize the effects of independent variables in non-linear models and they equal averages of the independent variables' effect over all the sample. We obtain

[18] Filipe, Grammatikos, and Michala (2016) shows similar evidence on a Portuguese dataset.

AMEs taking the empirical derivative of each logit with respect to each variable of interest.[19] To ease AMEs' interpretation we use standard deviation changes as the reference for all our balance sheet, credit and macroeconomic independent variables, while using one percent changes for the inflation and cost of borrowing series.

To explain with an example, let's focus on interest expenses over cash flow, the first variable in Figure 3 Panel a), presenting the AMEs of balance sheet variables. The colored bars describe how a one standard deviation increase in interest expenses over cash flow shifts the firms' probability of default *on average*, across all observations. The blue bar shows that the AME of a standard deviation increase in the ratio in year y on year y+1's default probability is slightly more than half a percentage point; the orange bar shows almost identical effects on year y+2; the grey on year y+3.

Concerning other firm-level variables in Figure 3 Panel a), we can see that the probability of default is lower for more capitalized/less leveraged firms, for those with more liquidity, and for those with greater profitability. On the contrary, a high level of interest expenses on cash flow and especially a high drawn over granted ratio (financial flexibility) contribute to increasing default risk.

In Panel b), recording economic sector dummies' AMEs above and beyond agriculture (the excluded category), we can observe that all sectors, but electricity and gas are more exposed to default than agriculture. Firms operating in manufacturing and construction are the riskiest, followed by transportation and retail trade.

Finally, in Panel c), we report the effects of changes in the macro variables, displaying an aggregate impact encompassing the effects from the different lags of each variable included in the forecasting equations.[20] The overall effect of nominal interest rates in each model is positive and especially large at the final forecast horizon of three years ahead. Instead, GDP growth, higher house prices and greater average profitability lower default rates, as expected. Finally, the statistical evidence suggests that inflation tends to be correlated with higher default risk, pointing to increased costs as the main effect of inflation on the average firm over the period covered by our estimation sample.

[19] We compute the AMEs by fitting our logit equation and obtaining an estimated default probability $\hat{p}$ for each firm. From the fact that $p = \exp(\mathrm{B}X)\ /(1 + \exp(\mathrm{B}X))$, where X is the full independent variable matrix in equation (1) and B its parameter matrix made of β elements, it follows that the partial derivative of p in a single variable equals to $\beta * p * (1 - p)$. Substituting the theoretical p with the average of the $\hat{p}$s we get the AME estimate. We perform these operations employing the margins package in R (see Leeper, 2017) and provide a deeper review on AME calculation and interpretation in Annex C.

[20] We refer to Section 4.2 for the list of lags included for each macroeconomic variable.

Figure 3 – Parameter estimates

a) Micro variables

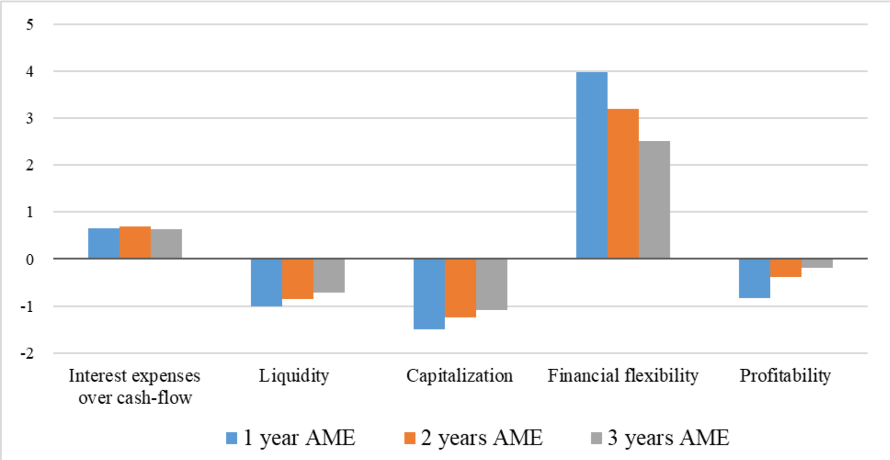

b) Economic sector

c) Macro variables

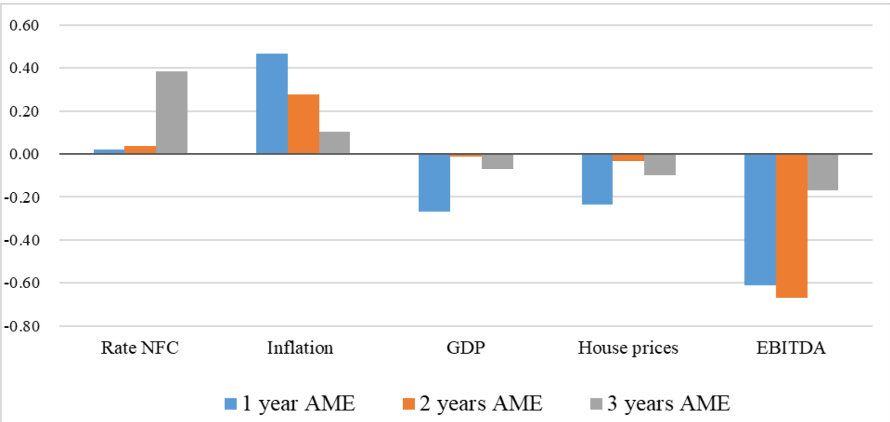

Source: Authors' calculation based on Cerved, Credit register data, Supervisory reports and Istat. *Note*: Across all panels, AME stands for Average Marginal Effects. In Panel b, we present the marginal effect of each sector with respect to a base sector (agriculture).

We move forward describing the heterogeneity in the effect of macroeconomic variables across sectors. We do so by obtaining the macroeconomic variables' AMEs averaging over observations pertaining to each specific industry alone. The non-linearity of the logit model implies that the final effect of each variable depends on all the effects of the other variables. As a consequence, local effect averages are meaningful impact heterogeneity's statistics.[21] The variability of impact estimates is crucial information, since banks often specialize in lending to a subset of firms (Paravisini, Rappoport, and Schnabl, 2022; Giometti and Pietrosanti, 2022; Duquerroy et al., 2022). As a result, a stark increase in one single sector's defaults may considerably harm specific banks and, through this, have unexpected systemic consequences.

We report the sectoral AMEs in Figure 4, on the y-axis, on a percentage points scale (e.g. 0,6 stands for 60 basis points). For example, the sub-plot showing the heterogeneity of inflation's effect across industries (Panel (a)) displays how an average one percent change in inflation (across all lags) affects the average default probability forecast for manufacturers, utilities, etc. In Panel (b), (c), and (d) we can instead observe the heterogeneity of GDP growth, NFC rate, and EBITDA's effects. We can see that the construction, accommodation, and transportation sectors are the most affected by macroeconomic shocks as of 2016's end.

The heterogeneity we report is economically significant. For just one example, looking at the effect of inflation on one-year-ahead defaults, the AME of a one percent increase in inflation for construction firms is an approximately 65 basis points increase in default probability. This local effect is 20 basis points larger than the baseline effect reported in Figure 3, the latter standing at about 45 basis points.

[21] An alternative would have been to estimate a fully interacted model. We did so and recorded very similar results in terms of ranking of sectoral exposures to different macroeconomic factors. Moreover, the fit and discriminatory capacity did not improve, while forecast error bands increased due to collinearity introduced by the coefficients' overabundance. We thus discarded the option. The cost of doing so is that we can only have more and less reactive sectors overall, at every point in time. For example, it cannot be that one sector is more reactive to GDP, and less to inflation. Nevertheless, being more or less reactive to macroeconomic shocks captures a key aspect of fragility; it is thus an important information to track.

Figure 4 – Sector specific reaction to changes in GDP and inflation

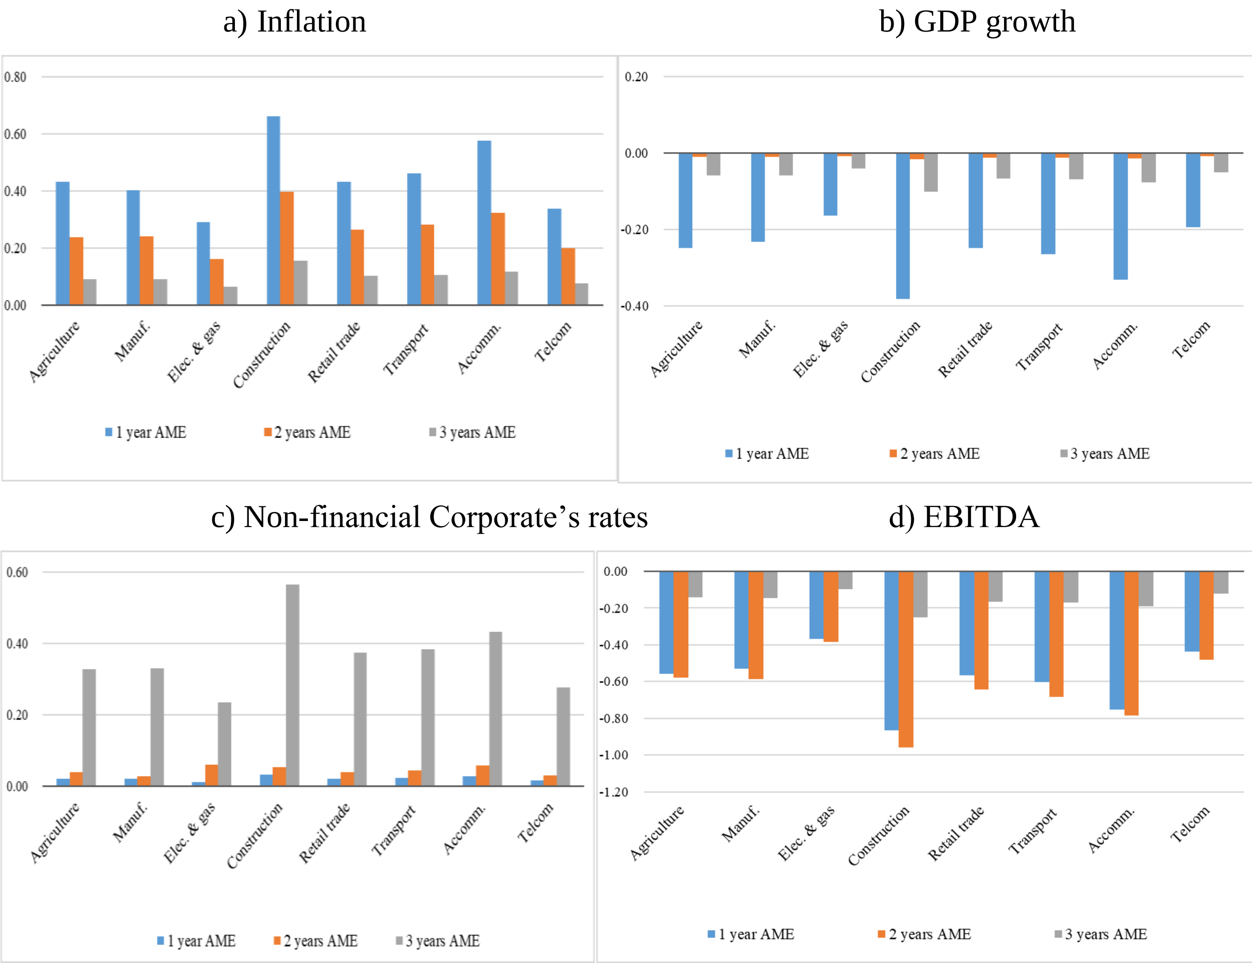

Source: Authors' calculation based on Cerved, Supervisory reports, Credit register data and Istat. *Note*: Across all panels, AME stands for Average Marginal Effects, displaying the sector-average effect of, respectively, one percent increase in inflation (Panel a), one standard deviation increase in GDP growth rate (Panel b; one s.d. change in GDP growth rate amounts to a 1,7 percent change), and one percent increase in NFC rates (Panel c).

In conclusion, jointly estimating a set of logit equations with macro and firm-level controls we find economically significant variability in the impact of macroeconomic factors on different industries. Heterogeneity's empirical relevance stresses the importance of having a flexible risk-analysis tool which directly embeds macro factors and can be aggregated at will.

4.5 Validation

We assess the prediction performance of the default models by considering two aspects. The first is the ability to fit the aggregate default rate, and the second is the power to discriminate defaulting and non-defaulting firms. Figure 5 plots the default rate over time, against the fitted average value. We present in separated panels the fit for the one, two, and three-years-ahead forecasts obtained with their respective model. The models replicate the default evolution over time with a very-good in-sample fit.

Figure 5 – All three equation have good in-sample performance

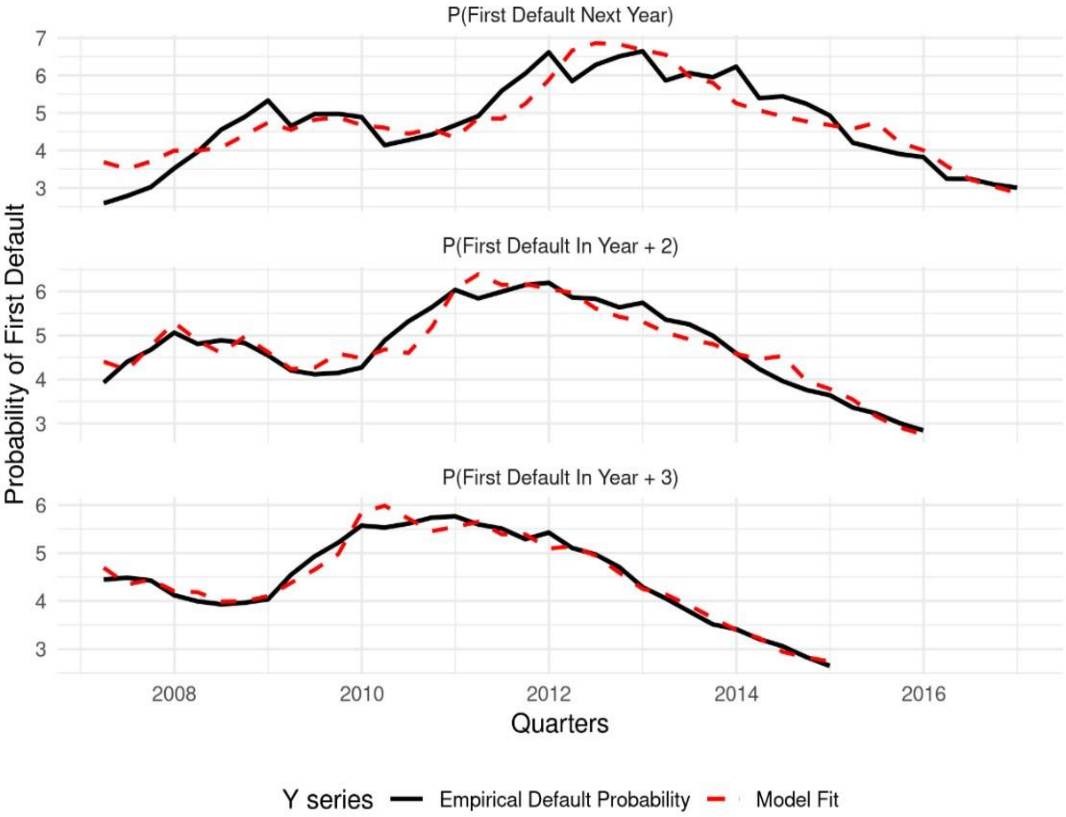

Source: Authors' calculation based on Cerved, Supervisory reports, Credit register data, and Istat. *Note*: The Figure presents model fits for the three forecast horizons we consider. For the one-year-ahead model, the estimation period is 2007Q1-2016Q4; 2007Q1-2015Q4 for the two-year-ahead model; 2007Q1-2014Q4 for the three-year-ahead one. This displacement allows us to recover the right parameters capturing the relation between ever more lagged firm level information. Assuming to have default data up to 2017 Q4, to obtain the one-year-ahead parameters we cut the last year of available information and forecast defaults between 2017 and 2018 Q4 with 2018 macro and 2017 firm level information. For the two-year-ahead model, we must cut two years, focusing on two-years-ahead defaults, which will only be available up to 2015Q4. Likewise, for the three-years-ahead model.

Second, we verify that discrimination capacity is good at all horizons. We report the results in Table 2. The full sample AUROC, calculated in-sample, is again close to 80 percent for each model. Figure 6 illustrates the distribution of the fitted PD one-year-ahead, splitting the plot between future defaulting and surviving firms. We can observe that the density is concentrated at zero for the non-defaulters, whereas the estimated probability is higher for firms that will default. This difference in densities supports the evidence of a good prediction performance at the borrower level.

Table 2: Default/Non-Default discrimination is good for all models, at all horizons

Model	Reference model
One-year-ahead	0.806
Two-year-ahead	0.772
Three-year-ahead	0.751

Source: Authors' calculation based on Cerved and Credit register data. *Note*: The Table compares power of discrimination between future defaulting and surviving firms for the reference model at the three forecast horizons. The reference model's number in the first cell differs from what shown in Table 1 since, there, we present the performance for the training set, here for the fully fledged model used in forecasting.

Figure 6: Firms classified as defaulters default more often

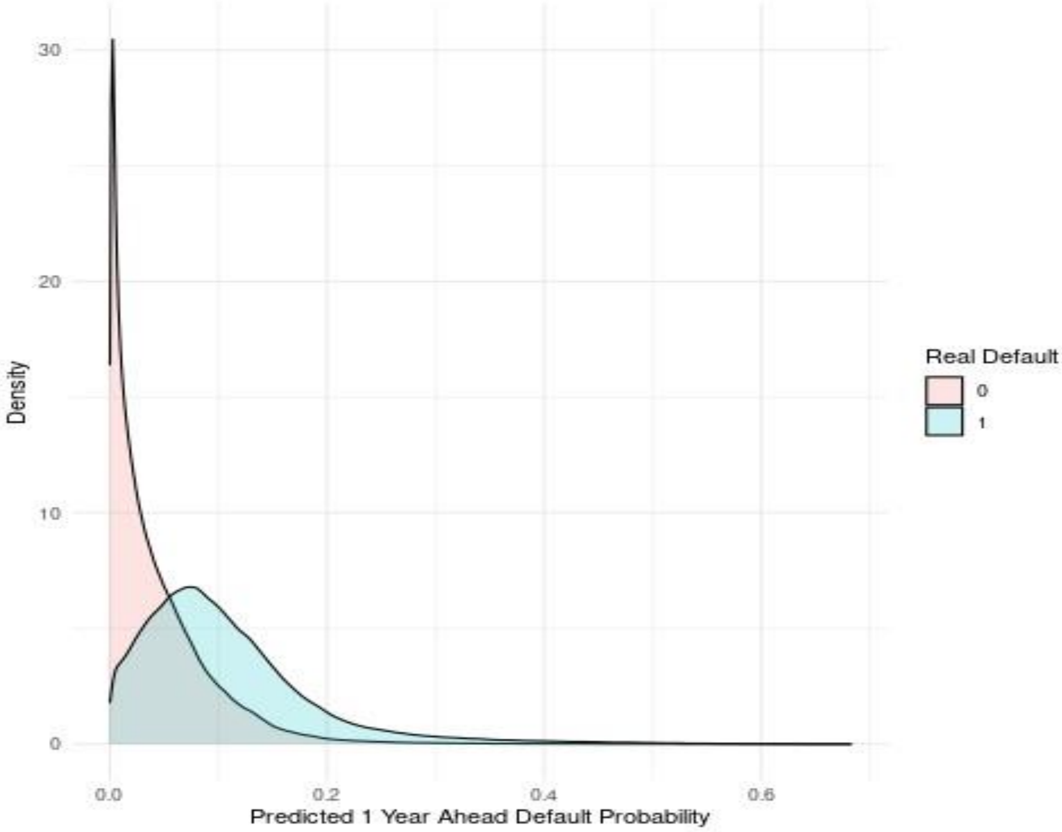

Source: Authors' calculation based on Cerved and Credit register data. *Note*: The Figure reports the density distribution for the estimated one-year-ahead probabilities of default distinguishing between firms that will default in one year (blue) and firms that will not (red).

5. Scenario Analysis

In this Section, we use our three-equation model to conduct simulation exercises, forecasting firms' default probability over a three-year horizon, starting in 2017, using realized macroeconomic time series as the baseline scenario and hypothetical pejorative realizations of the same time series as the adverse one. We start learning from our firm-level granular projections by aggregating them economy-wide in a forecast of the average default rate for corporates. This number is interesting per se, as an out-of-sample validation exercise. Furthermore, performing a full scenario analysis allows us to showcase useful applications of the at-will-aggregation potential of our model.

Starting from the firm-level projections, we explore the heterogeneity in default risk. As a first example, we present local averages of firm default probabilities at the industry and size-classes levels. Such local average forecasts allow us to explore the drivers of our economy-wide forecast. Moreover, through them, we check that our predictions align with our fundamental understanding of the correlation between risk and firm heterogeneity -- smaller firms and sectors like construction tend to be among the riskiest, and our results confirm such knowledge.

We then propose two last exercises that exploit forecast granularity. First, we look at the entire default probability forecasts' distribution for different industries and compare them under the same macroeconomic scenario.[22] Industries with similar average default rates could display very different tails. Averaging could hide pockets of vulnerability in some sectors that we may uncover thanks to our granular projections. Second, our model's forecasts naturally lend themselves to building transition matrices across risk classes when the conjuncture worsens. By tracking how many and which kind of firms transition from lower to higher risk classes when we change the scenario from baseline to adverse, we can learn which firms are more fragile to a change in the economic conjuncture.

5.1 The aggregate forecast exercise

We begin with a simple exercise, estimating the default models for one, two, and three years ahead. Each time, we forecast firm-level default probabilities based on the most recent balance sheet and macroeconomic time series describing a baseline and an adverse scenario. We derive the baseline scenario from the actual macroeconomic time series observed between the end of Q4 2016 and Q4 2019 (Bank of Italy elaborations based on ISTAT data). Instead, we build the adverse scenario by simulating changes in macroeconomic time series based on the three years of the European Sovereign Debt Crisis (2011 – 2013). Specifically, we assume equivalent changes from the Q4 2016 baseline for each macroeconomic time series over the years 2017-2019. (See Annex D for the macroeconomic variables under both the baseline and adverse scenarios, along with a brief scenario narrative). Finally, we calculate aggregate forecasts by averaging the firm-level default probability estimates across all firms.

In Figure 7, we present the GRANMA-PD model projections, from the end of Q1 2017 to the end of Q4 2020, shown against a light blue background. First, we observe that our one-, two-, and three-year-ahead out-of-sample projections, based on realized macroeconomic data and firm-level data from Q4 2016, accurately forecast the average default rate, validating our approach. Under the adverse scenario, the model predicts an average default rate just below 4 percent in 2018, compared to 3 percent at the end of 2017. Notably, the magnitude of this increase in the model predictions is in line with the rise in default rates observed in Italy during the peak of the Sovereign Debt Crisis. Indeed, between 2011 and 2012, the average default rate increased from approximately 5 percent to about 6.5 percent, representing a similar 30 percent rise.

The exercise period ends with the onset of the COVID-19 pandemic in Q2, Q3 and Q4 2020. During this last period, we interpret our forecast as projecting the average default path in the absence of the pandemic. This interpretation arises from the information set determining our 2020 forecasts, which stops with the end-Q4 2019 values of our macroeconomic time series, excluding the macroeconomic effects of the pandemic from our inputs. Despite this limitation, our baseline forecasts closely approximate the actual realization of the average default rate. We explain this result with two key

[22] The choice of comparing different industries has the only purpose of being an example. We could have chosen firm size, bank that provided the most credit, or any other relevant dimension as alternative markers of heterogeneity to monitor.

factors. First, in our dataset, a debtor enters the default state only when more than 5 percent of its outstanding credit is more than 90 days past due. As a result, the pandemic's effects on defaults could manifest only from Q3 2020 onward, as only then do the firms gravely affected by the March lockdowns start entering the default state. Second, several public support measures introduced in Q2 2020 helped to minimize deviations between the average default rate and its pre-pandemic trajectory, at least in the short term.

Figure 7: Projections based on the baseline and adverse macro scenarios

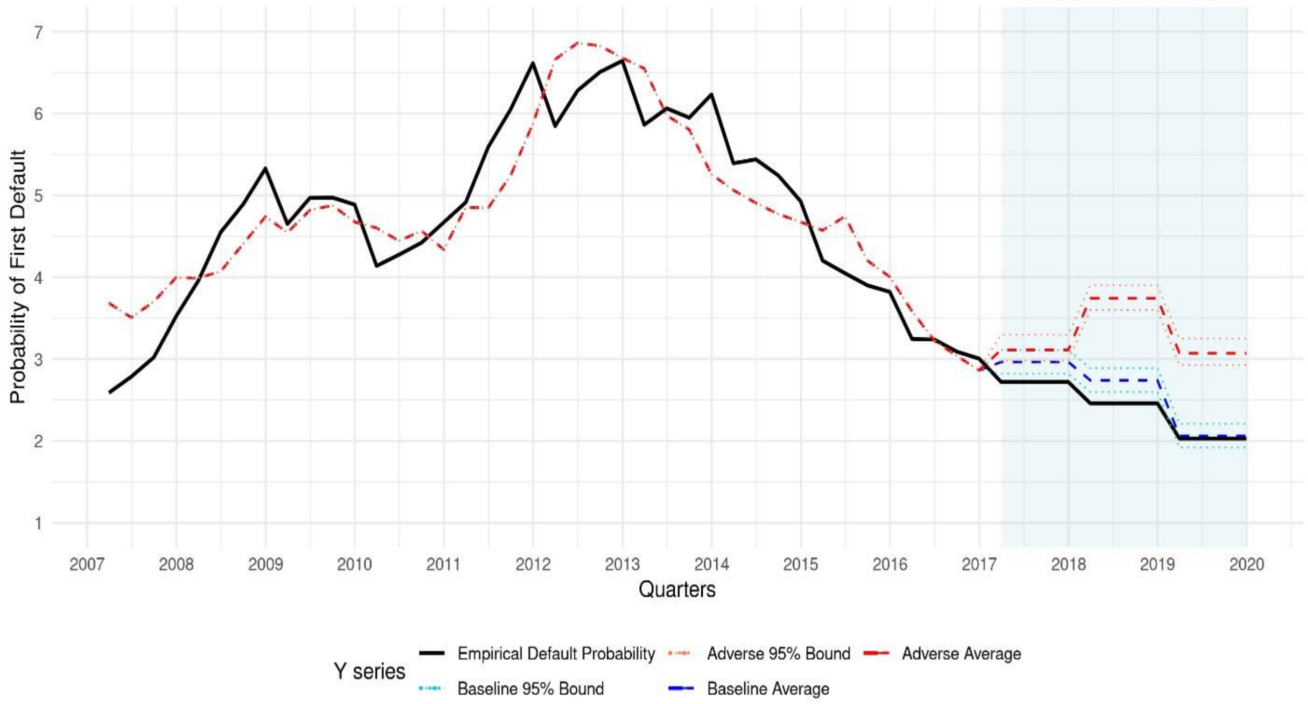

Source: Authors' calculation based on Cerved and Credit register data. *Note*: Estimation based on the three logit equations. X-ticks stand for last quarter of the year (e.g. 2017 means Q4 2017).

In Figure 8, we then exploit our model's granularity to examine the forecasts' heterogeneity across sectors and firm size classes.[23] We can see that the industry whose average default probability rises most in absolute terms is construction, with a 5.5 percent default rate peak under the adverse scenario, while all other sectors' forecasts under the same scenario hover between three and four percent. Furthermore, if we aggregate our forecasts over firms' size classes instead, we can see that micro and small firms primarily drive the peaks in the default forecast. In conclusion, our model not only reproduces the aggregate time series behavior of the average default rate, but it also captures cross-sectional regularities across different types of firms.

[23] As concerns the size classes, we use the EU Commission's definition of micro, small, medium and large firms.

Figure 8: Disaggregated projections

a) Industry heterogeneity

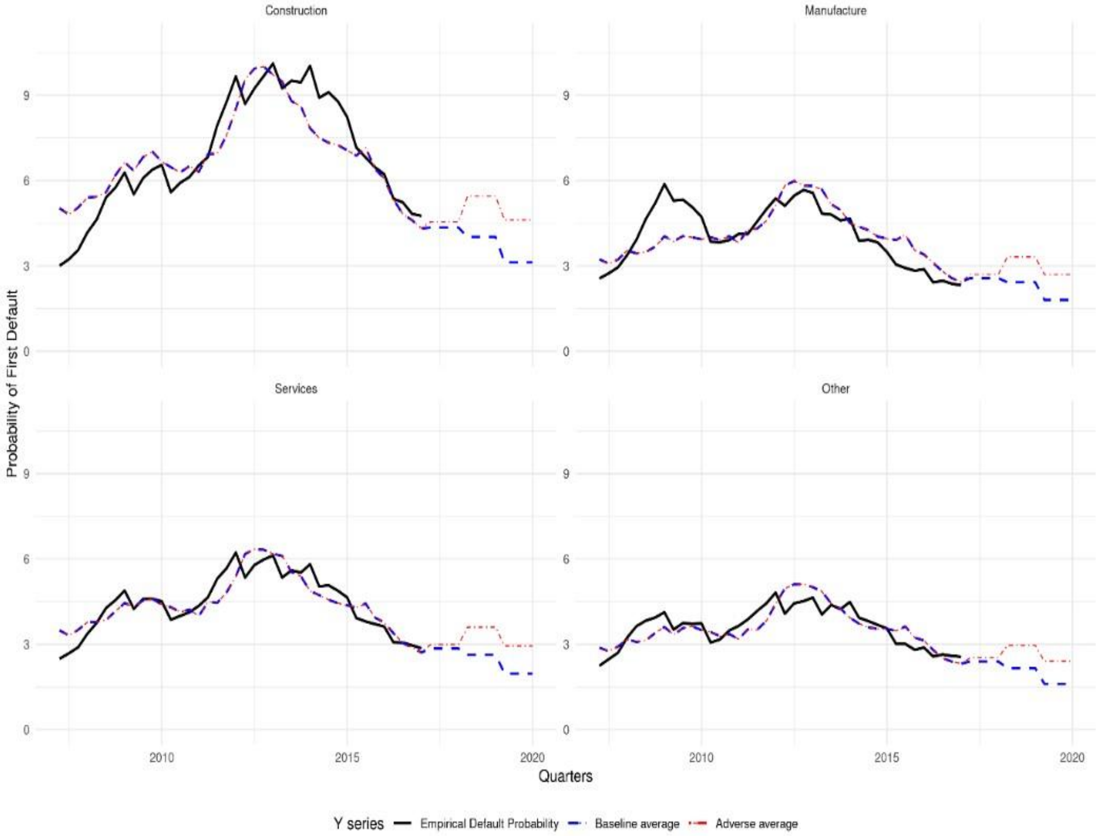

b) Size heterogeneity

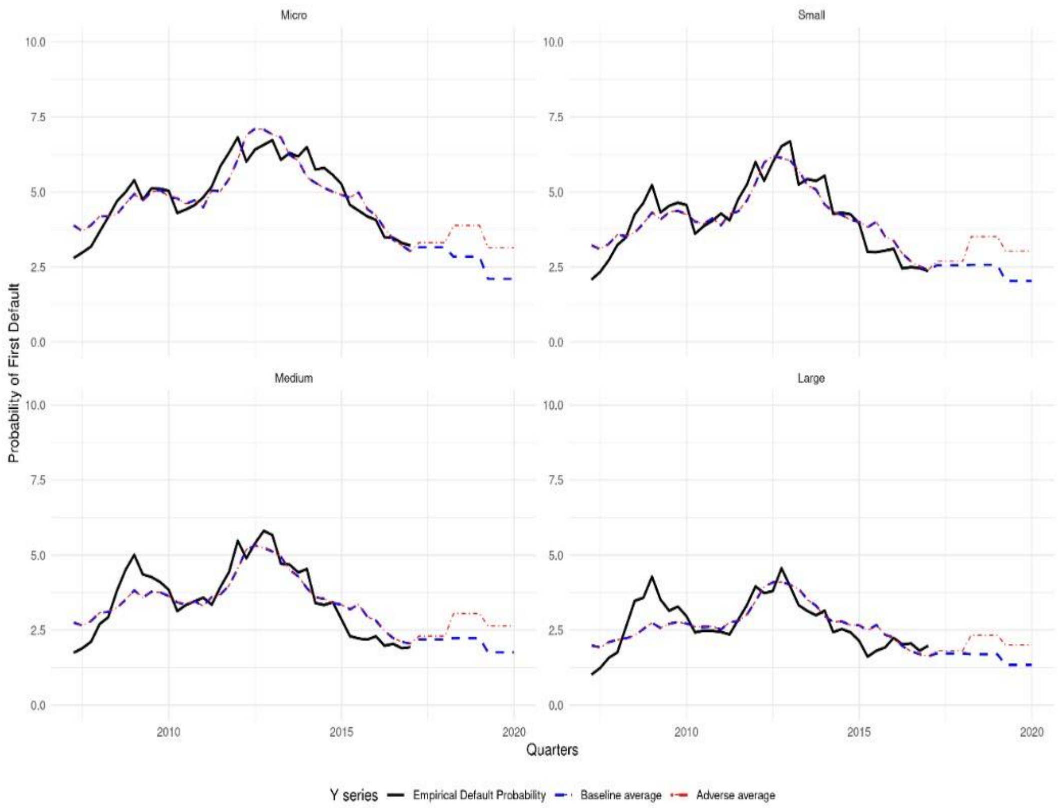

Source: Authors' calculation based on Cerved and Credit register data. *Note*: Estimation based on the three logit equations.

5.2 The granular forecasts exercise: Forecast distribution differences

We then exploit our model granularity more deeply, and use the forecasts to check for pockets of risk across different types of firms. We start with a within-scenario exercise meant to track divergence from the average in the distribution of the firm-level default probability forecasts for specific sectors. To provide an example, we display in Figure 9 the two-years-ahead PD distributions for the whole sample (pink) and specific industries (light blue).[24] In Panel a), we compare the average distribution to manufacturing's, and, in Panel b), to construction's. As expected, we observe that firms working in construction are more fragile, i.e., with a fatter-than-average right tail of the forecast distribution.

Figure 9: Sectoral risk indicators within the baseline scenario

a) Manufacturing b) Construction

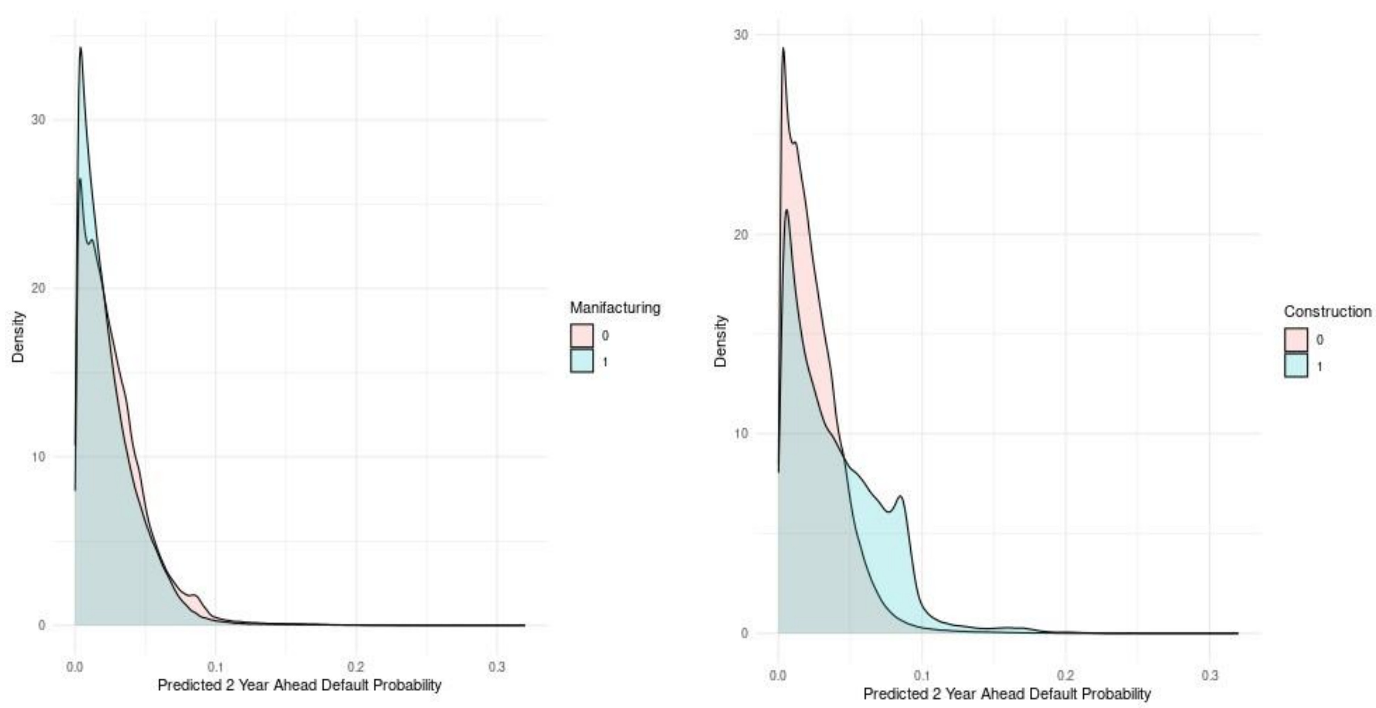

Source: Authors' calculation based on Cerved and Credit register data. *Note*: The Figure reports the density of the estimated defaults for the whole sample of firms (red) against those of firms belonging to a specific economic sector (blue). We estimate the probability of default by using the logit equation for defaults two-years-ahead.

5.3 The granular forecast exercise: Transitions between risk-classes

We propose an approach to assess the sensitivity to the change in the scenario of different subsets of firms, identified by their riskiness in terms of probability of default. First, we save the fitted default probability for each firm under the baseline and adverse-scenarios over the two-year forecast horizon; (again, the focus on two years ahead is just to provide an example). We then use the model predictions from the most recent year to determine the risk quartiles, thus defining four risk classes. Next, we observe which percentage of firms falls in each risk class in the baseline and adverse scenario and track transitions.

The resulting matrix reports the fraction of firms that migrate between quartiles due to the change in the scenario, keeping thresholds fixed, helping us tracking risks from adverse macroeconomic realizations. For example, in Figure 10, Panel A, we observe that 27 percent of firms were in the lowest risk class under the baseline scenario. Among these, 20 percent stay in the same risk class

[24] We choose the two-years-ahead horizon just to provide an example, we could have similarly chosen year plus one or year plus three, or displayed all of them.

under the adverse scenario, meaning that their default probability forecast under the adverse scenario is still below the threshold for class 1, and 7 percent see their probability of default rising enough to enter class 2. No firm shifts from class 1 to the higher risk classes (3 and 4). Overall, the transition matrix can help us trace granular vulnerability to sudden changes in the conjuncture.

Figure 10: Sectoral risk indicators between scenarios

A) Total firms

		Adverse				
		class 1	class 2	class 3	class 4	total
Base	class 1	20%	7%	0%	0%	27%
	class 2	0%	14%	13%	0%	27%
	class 3	0%	0%	11%	14%	25%
	class 4	0%	0%	0%	20%	20%
		20%	21%	24%	34%	

B) Construction sector

		Adverse				
		class 1	class 2	class 3	class 4	
Base	class 1	16%	6%	0%	0%	22%
	class 2	0%	10%	9%	0%	19%
	class 3	0%	0%	8%	11%	19%
	class 4	0%	0%	0%	40%	40%
		16%	16%	17%	51%	

C) Manufacturing sector

		Adverse				
		class 1	class 2	class 3	class 4	
Base	class 1	24%	9%	0%	0%	33%
	class 2	0%	15%	13%	0%	28%
	class 3	0%	0%	10%	12%	22%
	class 4	0%	0%	0%	16%	16%
		24%	24%	23%	28%	

Source: Authors' calculation based on Cerved and Credit register data. *Note*: We estimate PDs by using the logit equation for defaults two-years-ahead, and employing baseline (rows) and adverse scenarios' (columns) values for the macro variables.

6. Conclusions

This work presents a modelling framework relying on both micro and macro data for predicting corporate default risk under different macroeconomic scenarios. Our objective is to build a tool to forecast the probability of default at the firm level, aggregate these forecasts at will, and perform macro-scenario analysis taking granular firm characteristics into account. We apply the local projections' logic to logit hazard models. Instead of projecting one equation in the future, which would require forecasts of the microeconomic variables in addition to the macroeconomic ones, we estimate a model made of multiple equations, each local to a specific forecast horizon.

We base our estimates on a dataset covering 20 million observations referring to 840 thousand firms observed over 16 years. We build the model so that it has a good discrimination capacity between defaulters and non-defaulters while fitting the realized average default rate time series well. We further validate it with an out-of-sample forecasting exercise estimating the model using data up to 2017Q4 and projecting defaults for the years 2018 to 2020.

The average projections we display are aggregations of our model's granular forecasts. We show how such granularity can be an important value added per se by proposing two ways to learn more information on the evolution of risk from it. First, we show that comparing the forecasts' distributions across different types of firms (industries, but we could have used other groupings) can help us track the buildup of pockets of risk. Furthermore, our model's forecasts naturally lend themselves to building transition matrices across risk classes when the conjuncture worsens, helping us learn which type of firms bear more conjunctural risk.

References

Accornero, M., Alessandri, P., Carpinelli, L., & Sorrentino, A.M. (2017). Non-performing loans and the supply of bank credit: Evidence from Italy. *Bank of Italy Occasional Papers*, No. 374.

Altman, E.I. (1968). Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. *The Journal of Finance*, 23(4), 589-609.

Altman, E.I., & Hotchkiss, E. (2006). *Corporate Financial Distress and Bankruptcy*. 3rd Edition, John Wiley & Sons, Hoboken, NJ.

Altman, E. I., Marco, G., & Varetto, F. (1994). Corporate distress diagnosis: Comparisons using linear discriminant analysis and neural networks (the Italian experience). *Journal of Banking & Finance*, 18(3), 505-529.

Bacham D., & Zhao, J. (2017). Machine Learning: Challenges, lessons, and opportunities in credit risk modeling. *Moody's Analytics Risk Perspectives*, 9 (July).

Bakshi, G., Gao, X., & Zhong, Z. (2022). Decoding default risk: A review of modeling approaches, gindings, and estimation methods. *Annual Review of Financial Economics*.

Banerjee, R. N., & Blickle, K. S. (2016). Housing collateral and small firm activity in Europe. *Bank for International Settlements Working Papers*, No. 575.

Barboza, F., Kimura, H., & Altman, E. (2017). Machine learning models and bankruptcy prediction. *Expert Systems with Applications*, *83*(15), 405-417.

Bonaccorsi di Patti, E., & Cascarino, G. (2020). Modelling the dynamics of nonperforming loans in Italy. *Bank of Italy Notes on Financial Stability and Supervision*, No. 19.
Bruce, P., & Bruce, A. (2017). *Practical Statistics for Data Scientists*. O'Reilly Media.

Bühlmann, P. (2006). Boosting for high-dimensional linear models. *The Annals of Statistics*, *34*(2), 559-583.

Caldara, D., & Iacoviello, M. (2022). Measuring geopolitical risk. *American Economic Review*, *112*(4), 1194-1225.

Carling, K., T. Jacobson, J. Lindé, & Roszbach, K. (2007). Corporate credit risk modelling and the macroeconomy. *Journal of Banking & Finance*, *31*(3), 845–868.

Cascarino, G., Moscatelli, M., & Parlapiano, F. (2022). Explainable artificial intelligence: Interpreting default forecasting models based on Machine Learning. *Bank of Italy Occasional Paper*, No. 674.

Chava, S., & Jarrow, R.A. (2004). Bankruptcy prediction with industry effects. *Review of Finance*, *8*(4), 537-569.

Crouzet, N., & Mehrotra, N.R. (2020). Small and large firms over the business cycle. *American Economic Review*, *110*(11), 3549-3601.

De Socio, A., & Michelangeli, V. (2017). A model to assess the financial vulnerability of Italian firms. *Journal of Policy Modeling*, *39*, 147-168.

De Socio, A., Narizzano, S., Orlando, T., Parlapiano, F., Rodano, G., Sette, E. & Viggiano, G. (2020). Gli effetti della pandemia sul fabbisogno di liquidità, sul bilancio e sulla rischiosità delle imprese. Bank of Italy Covid-19 Note, November 2020.

Duquerroy, A., Mazet-Sonilhac, C., Mésonnier, J.S., & Paravisini, D. (2022). Bank Local Specialization. *Banque de France Working Papers*, No. 865.

Fabiani, A., Heineken, J., & Falasconi, L. (2022). Monetary policy and corporate debt maturity. *Available at SSRN 3945615*.

Filipe, S. F., Grammatikos, T., & Michala, D. (2016). Forecasting distress in European SME portfolios. *Journal of Banking & Finance*, *64*, 112-135.

Giacomelli, S., Mocetti, S., & Rodano, G. (2021). Fallimenti d'impresa in epoca Covid. Bank of Italy Covid-19 Note, January 2021.

Giesecke, K., Longstaff, F. A., Schaefer, S., & Strebulaev, I.A. (2014). Macroeconomic effects of corporate default crisis: A long-term perspective. *Journal of Financial Economics*, *111*(2), 297-310.

Giometti, M., & Pietrosanti, S. (2022). Bank Specialization and the Design of Loan Contracts. *Federal Deposit Insurance, Center for Financial Research Working Papers*, No. 14.

Giovannelli, F., Iannamorelli, A., Levy, A., & Orlandi, M. (2020). The in-house credit assessment system of Banca d'Italia. *Bank of Italy Occasional Paper*, No. 586.

Glennon, D., & Nigro, P. (2005). Measuring the default risk of small business loans: A survival analysis approach. *Journal of Money, Credit and Banking*, *37*(5), 923-947.

Gross, M., & Población, J. 2017. Implications of Model Uncertainty for Bank Stress Testing. *Journal of Financial Services Research*, Springer, 55(1), 31-58.

Ivashina, V., & Scharfstein, D. (2010). Bank lending during the financial crisis of 2008. *Journal of Financial Economics*, *97*(3), 319-338.

Jacobson, T., Lindé, J., & Roszbach, K. (2013). Firm default and aggregate fluctuations. *Journal of the European Economic Association*, *11*(4), 945-972.

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2014). *An Introduction to Statistical Learning: With Applications in R*. Springer Publishing Company, Incorporated.

Jensen, T. L., Lando, D., & Medhat, M. (2017). Cyclicality and firm size in private firm defaults. *International Journal of Central Banking*, *13*(4), 97-145.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, *95*(1), 161-182.

Jordà, Ò., Schularick, M., & Taylor, A.M. (2016). Sovereigns versus banks: Credit, crises, and consequences. *Journal of the European Economic Association*, *14*(1), 45-79.

Leeper, T. J. (2017). Interpreting regression results using average marginal effects with R's margins. *Available at the comprehensive R Archive Network (CRAN)*, 1-32.

Merton, R.C. (1974). On the pricing of corporate debt: The risk structure of interest rates. *The Journal of Finance*, *29*(2), 449-470.

Moscatelli, M., Parlapiano, F., Narizzano, S., & Viggiano, G. (2020). Corporate default forecasting with machine learning. *Expert Systems with Applications*, *161* (December), 113567.

Orlando, T., & Rodano, G. (2020). Firm undercapitalization in Italy: business crisis and survival before and after COVID-19, Bank of Italy Occasional papers No. 590, December 2020.

Orlando, T., & Rodano, G. (2022). The impact of Covid-19 on bankruptcies and market exits of Italian firms. Bank of Italy Covid-19 Note, January 2022.

Ottonello, P., & Winberry, T. (2020). Financial heterogeneity and the investment channel of monetary policy. *Econometrica*, *88*(6), 2473-2502.

Paravisini, D., Rappoport, V., & Schnabl, P. (2022). Specialization in bank lending: Evidence from exporting firms. *The Journal of Finance*, *Forthcoming*. *Available at SSRN 3993978*.

Reinhart, C. M., & Rogoff, K.S. (2009). The aftermath of financial crises. *American Economic Review*, *99*(2), 466-72.

Schularick, M., & Taylor, A.M. (2012). Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870-2008." *American Economic Review*, *102*(2), 1029-61.

Shumway, T. (2001). Forecasting bankruptcy more accurately: A simple hazard model. *The Journal of Business*, *74*(1), 101-124.

——— (2019): “Financial Stability Review”, n.2, Bank of Italy.

ANNEX A: Model variables

Table A.1: Macro indicators

Variable	Description	Frequency	Source
GDP	GDP growth rate (seasonally adjusted)	Quarterly	Bank of Italy's elaboration on Istat data, Bank of Italy projections
Rate NFC	Average nominal interest rate on new loans to nonfinancial firms	Quarterly	Supervisory Report, Bank of Italy projections
Inflation	Growth rate of deflator for national final consumption (seasonally adjusted)	Quarterly	Bank of Italy's elaboration on Istat data, Bank of Italy projections
House prices	Average price per square meter of residential real estate; growth rate (seasonally adjusted)	Quarterly	Bank of Italy's elaboration on Istat data, Bank of Italy projections
EBITDA/Gross operating margin	Gross operating margin of nonfinancial firms (; index (base 100= 2008)	Quarterly	Bank of Italy's elaboration on Istat data, Bank of Italy projections

Table A.2: Firm level indicators

Indicator	Description	Frequency	Source
Capitalization	Ratio between equity and total assets. Used to assess how relevant are equity resources to fund a company's commitments.	Annual	Cerved
Liquidity	Ratio between current assets and short-term debt (current ratio). It measures each firm's ability to pay off short-term debt obligations with cash and cash equivalents.	Annual	Cerved
Firm size	Four classes according to the following criteria defined by European Commission: Micro, less than 10 employees and turnover (or total assets) up to EUR 2 million; small, less than 50 employees and turnover (or total assets) not exceeding EUR 10 million; medium, less than 250 employees and turnover less than EUR 50 million (or total assets not exceeding EUR 43 million); large, all remaining firms.	Annual	Cerved
Interest expenses to cash flow	Ratio that indicates the enterprise's ability to pay interest from generated cash flow.	Annual	Cerved
Profitability	Ratio between gross operating margin and total assets. It measures how much earnings the firm is generating before interest, taxes, depreciation, and amortization, as a percentage of total assets.	Annual	Cerved
Financial flexibility ratio	The ratio between drawn amount and granted amount. It measures the percentage of available credit that the firm is actually using. It refers to all the different types of loans.	Monthly	Central Credit Register
ATECO section	Dummy variables identifying firms' economic sector. We consider NACE code.		

Table A.3: Summary Statistics for Firm Level Indicators

Variable	N	Mean	St. dev.	Min	p 25	p 75	Max
Capitalization	15,388,258	0.200	0.229	-0.503	0.048	0.313	0.888
Liquidity	15,388,258	0.067	0.109	0	0.003	0.081	0.625
Interest expenses to cash flow	15,388,258	0.957	1.302	0	0.136	1.667	1.957
Profitability	15,388,258	0.066	0.125	-0.444	0.014	0.118	0.486
Financial flexibility	15,388,258	0.723	0.296	0.001	0.525	0.996	1.094

Source: Cerved and Credit Register. *Note*: See Table A.2 for definitions.

ANNEX B – Bootstrapped standard errors

Estimates' volatility evaluation is one of the goal of the analysis, since we want to understand what could be the deviation from the provided point estimates in a reasonable probability range. In order to address the uncertainty associated with prediction errors and conditional means, we need to build a proper prediction interval. Considering the modelling framework in which we move (generalized linear models), the standard calculation used for OLS regression based on Gaussian distributional assumption is not suitable. Estimating Equation (1), we cannot simply add and subtract twice the estimated standard errors to get the needed interval. We follow two methods to tackle this problem: endpoint transformation and bootstrapping.

The **endpoint transformation method** computes confidence intervals for monotonic functions of $x'\beta$, such as the predicted probabilities in binary logit. It starts from computing the symmetric confidence interval for the linear predictor ($x'\beta$). After that, it just applies the logit transformation on the upper and lower bounds. This method has the advantage of being computationally less intensive than bootstrapping, while ensuring that the final bounds of the prediction interval cannot exceed the range [0,1]. Nevertheless, it is a pure mathematical transformation that follows a simple rule: If the predicted probability is lower than 50 per cent (being always the case in this work) the lower bound will be always closer to the point estimate than the upper bound, and vice-versa. As such, the asymmetric distance between estimates and confidence bound will carry no economic meaning.

Bootstrapping procedure overcomes this limitation, since its asymmetrical prediction intervals are not predetermined by any mathematical rule. The intuition is that we can measure actual population sampling variability by taking random samples from the original dataset. In detail, the procedure works as follow:

1. Extract a random sample from the original dataset taking observations with replacements (the dataset size remains the same);
2. Repeat Step 1 as many times (e.g. 1,000) as the number of desired models to re-fit;
3. Re-fit the model on the samples extracted;
4. Generate the new estimates for every repetition (so every observation in the dataset will have 1,000 point estimates);
5. Collect all the estimates and calculate the needed percentiles (if you want a 95 per cent interval, take the 2.5 per cent and the 97.5 per cent percentiles as lower and upper bound respectively)

Bootstrapping fully exploits the informative power hidden in the data. The downside, though, lies in its computational cost: Repeated sampling and estimation, given our large dataset size, can extremely slow (days to obtain estimates) or unfeasible. For this reason, we apply a shortcut: We verify that taking a random 10 percent subsample does not affect materially our logit point estimates or prediction. Having verified this, we obtain the point estimates from the full sample, while the confidence intervals from the 10 percent subsample. In the future, increasing the computational power at our disposal, we plan to broaden this computation to the full dataset. Still, we do not expect significant changes in the results.

ANNEX C – Marginal effects in logistic regressions

The interpretation of logit model parameters in terms of predicted probabilities is not straightforward. To study the contribution of each variable on the final outcome, we need thus to calculate the **marginal effects (ME)**. In a regression framework, they are defined as the effect on the target variable of a change in the explanatory variable under investigation.

From a mathematical perspective, they are computed taking the derivative of the regression function with respect to the explanatory variable of interest. This is an easy task in the standard linear regression:

$$E(y \mid x_1, \dots, x_k) = \beta_0 + \beta_1 x_1 + \cdots + \beta_k x_k$$

The ME of x_1 is:

$$\frac{\partial E(y|x_1, \dots, x_k)}{\partial x_1} = \beta_1$$

The calculation is the same for every variable. So, we can draw two conclusions: in linear regression marginal effects overlap with estimated parameters, plus they are constant for every observation in the sample.

Things are more complicated in a non-linear framework, as for logistic regression. In this case, the expected value of the dependent variable, conditional on the independent ones, is:

$$E(y \mid x_1, \dots, x_k) = F(\beta_0 + \beta_1 x_1 + \cdots + \beta_k x_k)$$

where F(-) is the logistic function. The ME of x_1(**provided it is a continuous variable**) is:

$$\frac{\partial E(y|x_1,\dots,x_k)}{\partial x_1} = \beta_1 \text{ F'}(\beta_0 + \beta_1 x_1 + \cdots + \beta_k x_k)$$

where F'(-) is the probability density function of the logistic distribution. We can immediately appreciate that:

1. We cannot interpret the estimated parameter β_1 as ME anymore, as done for linear regressions.
2. The ME of x_1 **is different for every observation in the sample**, since it is a function of the values of all the covariates and of the estimated parameters (this makes sense, given the non-linearity of the logistic regression).

To evaluate the rate of change of the predicted probabilities resulting from a small[25] change in the continuous explanatory variables of interest, we introduce the Average Marginal Effect (AME). We calculate the ME per each observation in the sample (exploiting the latter equation above) and then we take the mean. In formula:

$$AME(x_1) = \beta_1 \sum_{i=1}^{n} \frac{\text{F'}(\beta_0 + \beta_1 x_{1,i} + \cdots + \beta_k x_{k,i})}{n}$$

The interpretation is the following: on average, a small change of $x_1$is associated with a change of AME(x_1) percentage points in the probability of y occurring.

The ME calculation is different for the **categorical explanatory variables**, since it doesn't make sense to refer to a 'small' change in their values. In this case the AME is equal to the change in the predicted outcome when the variable is set to a given level (i.e. Construction) minus the predicted outcome when the variable is set to its baseline level. To clarify, let's assume we have only two sectors of economic activities: Construction and Agriculture (baseline). We compute the AME for Construction as follow:

[25] $(x_1 + h)$, having h→0

1. For the first observation in the sample, fit the predicted probability through the model built, forcing the economic activity sector to be Construction and leaving the other independent variables as they are.
2. Repeat step 1, but in this case force the economic activity sector to be Agriculture.
3. the ME for the individual case is the difference in predicted probabilities from step 1 and 2;
4. repeat the process for every observation in the sample and take the mean: this is the AME for Construction.

ANNEX D – Macroeconomic variables scenarios

Baseline scenario. Over the three years, the Italian economy experiences a weak economic growth and a reduction in inflation; meanwhile, interest rates have been decreasing, improving debt sustainability and helping containing the rise of macroeconomic risks. After a prolonged period of house price decline, the real estate cycle remains weak, even though some signs of recovery materialize toward the end of 2019.

Adverse scenario. Over the first two years, the Italian economy experiences a strong negative GDP shock, lowering firms' profitability significantly. While inflations stays low and keeps decreasing, the severely worsened real conjuncture implies a sharp rise in the perceived risk by intermediaries, whose supply tightens, as reflected in a 30 percent increase in the interest rate faced on average by non-financial corporates. Over the last year, the Italian economy stops contracting, but does not grow either.