GENEVA GRADUATE INSTITUTE

INSTITUT DE HAUTES ÉTUDES INTERNATIONALES ET DU DÉVELOPPEMENT

GRADUATE INSTITUTE OF INTERNATIONAL AND DEVELOPMENT STUDIES

Graduate Institute of International and Development Studies

International Economics Department

Working Paper Series

Working Paper No. HEIDWP29-2026

Financial Stress Indicator for Peruvian Non-Financial Corporations

Gerald Cisneros
Central Reserve Bank of Peru

Marcos Cerón
Central Reserve Bank of Peru

Jose Apaza
Central Reserve Bank of Peru

Chemin Eugène-Rigot 2
P.O. Box 136
CH - 1211 Geneva 21
Switzerland

©The Authors. All rights reserved. Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. No part of this paper may be reproduced without the permission of the authors.

Bilateral Assistance & Capacity Building for Central Banks

Financial Stress Indicator for Peruvian Non-Financial Corporations

Gerald Cisneros
Marcos Cerón
Jose Apaza

Central Bank of Peru (BCRP)

Abstract

This paper develops a financial stress indicator for Peruvian non-financial corporations using annual financial statements over the period 2000–2025. First, we evaluate the gains in predictive performance from recalibrating the original Altman model using Peruvian data. We then construct a broader set of financial indicators from six dimensions: profitability, liquidity, leverage, coverage, activity, and growth and size, to estimate a range of distress models, including static logit, panel logit, and hazard specifications. The results show that the original Altman model performs poorly when applied directly to Peruvian firms, and the recalibration improves predictive accuracy. Expanding the information set beyond the Altman ratios yields further gains in terms of accuracy and reductions in both Type I and Type II classification errors. Using the estimated firm-level probabilities, we construct measures of corporate financial stress based on both asset- and debt-weighted aggregation schemes. The resulting indicators successfully track major episodes of corporate vulnerability over the sample period, including the global financial crisis and the COVID-19 pandemic. A Shapley decomposition reveals that profitability is the main dimension associated with corporate distress.

Keywords: financial distress; bankruptcy prediction; financial stress indicator; LDA, logistic regression, LASSO; emerging markets; Peru.

JEL: G33, G32, C52, C53, E44.

We are especially grateful to Katja Neugebauer for her continuous guidance, generous time, and detailed feedback as our coach throughout the development of this paper. We also gratefully acknowledge the support of the Geneva Graduate Institute through the BCC programme, and thank Cédric Tille and seminar participants for insightful comments and suggestions that helped improve the paper.

The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Central Reserve Bank of Peru. All remaining errors are our own.

1 Introduction

The corporate sector is one of the largest counterparties of the Peruvian financial system. Non-financial corporations account for around 35% of total bank loans and are the main issuers in the domestic fixed-income market. Deterioration in corporate balance sheets can therefore affect financial stability through several channels, including loan losses, changes in the valuation of corporate debt, and collateral positions. In consequence, monitoring the financial condition of non-financial corporations is an important component of financial stability surveillance.

Despite this importance, there is no systematic indicator of financial distress for the Peruvian non-financial corporate sector. Existing approaches typically focus on individual firms or on specific financial indicators. This paper develops a framework that addresses this gap in the literature, by combining firm-level distress prediction and aggregate financial stress indicator.

The starting point is the Altman (1968) Z-score, one of the most used distress metrics among rating agencies and supervisors. However, applying the model with its original coefficients and cutoff to Peruvian firms may not be the best approach. In a previous study, using a thirty-one country sample and more than six million observations, Altman et al. (2017) found that no single financial distress model can work globally: coefficients and thresholds must be recalibrated to the country and its firm setting. Considering this, the first exercise compares the 1968 model against a locally recalibrated model on a common test sample, measuring the accuracy gains from recalibration.

Then, we ask if the predictive performance can be improved by moving beyond the five ratios used in the original Altman model. We organise a broader set of financial indicators into six economic dimensions: profitability, liquidity, leverage, coverage, activity, and growth and size. This broader information set allows us to evaluate whether additional financial indicators contain predictive information that is not captured by the traditional Altman ratios. With this new set of predictors, we estimate additional models, including static logit, and dynamics specifications. Incorporating a broader set of predictors, however, raises a variable-selection problem. Financial ratios are highly correlated and may contain overlapping information, which makes the model vulnerable to redundancy and overfitting.

Earlier distress models relied on a small number of ratios selected using statistical significance, correlation, or prior evidence. More recent studies have instead employed penalised regression and machine-learning techniques to address the dimensionality of financial information (Tian et al., 2015; Vu et al., 2023; Qian et al., 2022; Papikova and Papik, 2022). Following this literature, we use a multi-stage selection procedure that combines univariate discriminatory screening, within-category redundancy reduction, and global de-correlation, which we compare with a LASSO-based selection. The resulting

parsimonious specification retains indicators from all six dimensions while substantially reducing the number of predictors.

A further question is whether incorporating the time dimension and macro-financial conditions improves the prediction of corporate distress. Static models treat firm-year observations primarily as cross-sectional information, whereas financial distress is a dynamic process and may depend on the length of time firms remain exposed to adverse conditions. Following Chava and Jarrow (2004) and Bharath and Shumway (2008), we compare the static specifications with panel-logit and discrete-time hazard models that incorporate macro-financial variables. All specifications are evaluated over the same out-of-sample period, 2020–2025, allowing their predictive performance to be compared under a common set of conditions that includes the COVID-19 shock.

The final step is to transform the firm-level probabilities generated by the model into an aggregate measure of corporate financial stress. Following the stress-index literature (Galán, 2021), we construct both an extensive margin, based on the share of firms classified as being in stress, and intensive margins based on asset- and liability-weighted predicted distress probabilities. This allows the indicator to capture whether distress is concentrated among firms with greater economic size or larger liability exposure.

Finally, to answer which financial dimensions account for changes in predicted corporate stress, we employed the Shapley decomposition to decompose the predicted pprobabilities

The analysis produces four main findings. First, recalibrating the original Altman model improves discriminatory power. On the common 2020–2025 test sample, the original model achieves an AUROC of 0.78, compared with 0.94–0.95 for the recalibrated specifications. Second, expanding the predictor set beyond the five Altman ratios produces further gains: the parsimonious logit model achieves an out-of-sample AUROC of 0.97, with substantial reductions in both Type I and Type II classification errors. Third, incorporating macro-financial variables through panel-logit and hazard specifications improves the detection of distress events but does not improve overall discriminatory power. Fourth, the resulting aggregate indicators capture major episodes of corporate vulnerability over the sample period, while the Shapley decomposition identifies profitability as the main financial dimension .

The paper contributes to the literature in three ways. First, it provides a evaluation of financial-distress models for Peruvian listed non-financial corporations using more than 20 years of financial statements. Second, it develops a variable-selection framework that improves predictive performance without relying on a large set of ratios. Third, it links firm-level distress prediction to an aggregate financial stress indicator and provides a category-level decomposition that helps to identify the main drivers of corporate distress.

Related Literature

One of the earliest papers related to financial distress was Beaver (1966), who demonstrated in a matched-sample framework that individual accounting ratios can distinguish failing from non-failing firms before the actual failure. Later, Altman (1968) extended the analysis to a multivariate setting through the linear discriminant framework. From an initial set of 22 accounting and market-based variables, Altman selected five ratios representing liquidity, profitability, operating performance, market leverage, and asset turnover. These variables were combined into a single score with weights estimated by the linear discriminant analysis (LDA). Altman then estimated classification thresholds of 1.81 and 2.99, defining a region of uncertainty between distressed and non-distressed firms. Subsequent research broadened both the set of predictors and the statistical techniques employed. Altman et al. (1977) developed the ZETA model, incorporating additional dimensions such as earnings stability and debt-servicing capacity.

A different approach was proposed by Ohlson (1980), who replaced the LDA model with a logistic model estimated on a large unmatched sample of firms. Zmijewski (1984) employed a probit framework and emphasised the potential biases that arise when failure cases are deliberately oversampled relative to their population frequency.

With respect to the categories of potential predictors, Taffler (1983) argued that financial ratios reflect a limited number of underlying economic dimensions, including profitability, working-capital position, financial risk, and liquidity. Pinches et al. (1973) and Chen and Shimerda (1981) reached similar conclusions using factor-analytic methods, finding that a large set of accounting ratios collapses into a small number of stable dimensions.

A major critique of the bankruptcy-prediction literature was provided by Shumway (2001), who showed that traditional models estimated on a single observation per firm are misspecified because they ignore the length of time that firms remain at risk. Under a hazard model, he incorporates the dynamic dimension into the model and exploits the panel structure of the data, finding that several predictors lose statistical significance under this specification.

Subsequent studies explore the advantages of incorporating a dynamic framework. Chava and Jarrow (2004) extended Shumway's specification by incorporating industry effects, while Bharath and Shumway (2008) found that hazard models combining accounting and market information outperform the Merton distance-to-default framework in out-of-sample prediction. In a more recent application, Gupta et al. (2018) compared several models, including logit, panel logit and hazard specifications, for SME failure prediction and reported only minor differences in discriminatory performance.

In the case of the Peruvian market, Lizarzaburu (2014) evaluates the Altman Z'-score model on the Selective Index of the Lima Stock Exchange (ISBVL) from 2008 to 2012,

finding a 78% predictive effectiveness across representative firms.

2 Data

2.1 The definition of distress

Financial distress has no single definition in the literature. Wruck (1990) defines financial distress as a state in which a firm's cash flow is insufficient to meet its current obligations. Most studies employ different definitions depending on data availability or the objectives of the analysis, with most definitions falling into three groups. Because the definition of financial distress determines the event rate, the composition of the distressed group, the methodology and the empirical results, the choice is an important step in the study.

The first group is the legal-event definitions. Beaver (1966) classifies a firm as failed when it goes bankrupt, defaults on a bond, overdraws its bank account, or omits a preferred dividend. Altman (1968), Ohlson (1980), and Zmijewski (1984) use the filing of a bankruptcy petition as the relevant event. Shumway (2001) extends the event set to include performance-related delistings alongside bankruptcy filings. The main advantage of this type of definition is that the event is unambiguous and precisely dated. Its main disadvantage is that it requires distressed firms to enter formal proceedings and to be recorded in the data. Moreover, insolvency procedures can be slow and costly, and the resulting number of observed events may be small. In our dataset, formal bankruptcy and liquidation proceedings account for only 10 firm-year observations, which is insufficient for reliable forecasting.

The second group consists of market-based definitions, which rely on security prices and market-value variables. Campbell et al. (2008) define failure as bankruptcy, performance-related delisting, or a D credit rating. Bharath and Shumway (2008) use the Merton distance-to-default as a measure of financial distress, while Duffie et al. (2007) model rated defaults directly using market information and security prices. In the Peruvian stock market, however, most issuers trade infrequently, resulting in stale equity prices. In addition, corporate credit ratings cover only a small subset of issuers. These limitations make market-based definitions difficult to implement in our setting.

The third group defines financial distress on the basis of financial ratios or accounting information. Whitaker (1999) defines a period of financial distress as beginning in the first year in which cash flow falls below long-term debt. Asquith et al. (1994) classify a firm as distressed when EBITDA falls below interest expense for two consecutive years, or below 80% of interest expense in any single year.

Pindado et al. (2008) use a definition based on EBITDA falling below financial expenses for two consecutive years, combined with a decline in market value. Manzaneque et al. (2016), working with Spanish firms, relax the joint requirement by treating the two

conditions as alternatives, such that the coverage condition alone is sufficient to classify a firm as distressed.

Table 1: Definitions of financial distress in the literature

Family	**Explicit criterion**	**Window**	**Studies**
Legal event	Bankruptcy, bond default, bank overdraft, or omission of a preferred dividend	Event	Beaver (1966)
	Bankruptcy petition	Event	Altman (1968); Ohlson (1980); Zmijewski (1984)
	Bankruptcy filing or performance-related delisting	Event	Shumway (2001)
Market-based	Bankruptcy, performance-related delisting, or a D credit rating	Event	Campbell et al. (2008)
	Merton distance-to-default as a measure of financial distress	Continuous	Bharath and Shumway (2008)
	Rated default based on market information and security prices	Event	Duffie et al. (2007)
Accounting-based	Cash flow falls below long-term debt	1 year	Whitaker (1999)
	EBITDA < interest expense in two consecutive years, or EBITDA < 80% of interest expense in one year	2 years	Asquith et al. (1994)
	EBITDA < financial expenses in two consecutive years, combined with a decline in market value	2 years	Pindado et al. (2008)
	EBITDA < financial expenses in two consecutive years, with the coverage condition used as an alternative to the market-value condition	2 years	Manzaneque et al. (2016)
	EBIT < financial expenses in two consecutive years	2 years	Pham Vo Ninh et al. (2018)
	Negative net income in two consecutive years	2 years	Geng et al. (2015)

Notes. 'Window' refers to the persistence requirement: the number of consecutive years over which the condition must hold. 'Event' denotes a discrete legal or market occurrence with no persistence requirement. 'Continuous' refers to a continuous measure of financial distress rather than a discrete classification.

Since financial statements are available for the whole sample, the third group of definitions is the one we can implement. We classify a firm as financially distressed if its coverage ratio stays below one for two consecutive years, or if it enters formal insolvency proceedings. EBITDA would be the preferred measure of operating earnings, but it is not reported in our data, so we use EBIT as a proxy.

$$D_{it} = \begin{cases} 1, & \text{if Coverage}_{it} < 1 \text{ and Coverage}_{i,t-1} < 1, \\ 1, & \text{if FormalEvent}_{it} = 1, \\ 0, & \text{otherwise}, \end{cases} \tag{1}$$

where the coverage ratio is defined as

$$\text{Coverage}_{it} = \frac{\text{EBIT}_{it}}{\text{FinExp}_{it}}. \tag{2}$$

D_{it} is a binary indicator equal to one if firm i is classified as financially distressed in year t, and zero otherwise. FormalEvent_{it} is an indicator equal to one if firm i enters formal bankruptcy or liquidation proceedings during year t, and zero otherwise. EBIT_{it} denotes earnings before interest and taxes, while FinExp_{it} denotes financial expenses.

2.2 Sample and coverage

The main data source is the set of annual individual financial statements reported to the capital market supervisor by listed issuers between 2000 and 2025 (26 years).

Given that not all firms are present in every year of the sample, we divide the data into five mutually exclusive subsamples according to the number of years with reported information: the balanced panel (all years observed) and permanence bands of 15–25, 10–14, 5–9, and ≤ 4 years. Table 2 reports each subsample together with its distress rate.

Table 2: Sample by permanence in the panel

Subsample	Firms	Firm-years	Events	Distress rate
S0. Balanced panel	51	1,280	176	13.8%
S1. Permanence 15–25 years	67	1,216	242	19.9%
S2. Permanence 10–14 years	33	354	85	24.0%
S3. Permanence 5–9 years	55	345	82	23.8%
S4. Permanence ≤ 4 years	85	194	50	25.8%
Reference estimation sample (S0–S3)	206	3,195	585	18.31%

Notes: permanence is the number of years with reported data, and the subsamples are mutually exclusive. Source: authors' calculations.

In order to limit the effects of outliers, data-quality problems and firms with insufficient observations, our estimation sample consists of firms with at least five years of data (subsamples S0–S3). This yields 206 firms, 3,195 firm-year observations, and 585 distress events (18.31% of observations).

The subsamples also reveal a survivorship pattern. Distress rates are increasing as permanence declines, rising from 13.8% in the balanced panel to 23.8% among firms observed for five to nine years and 25.8% among firms observed for four years or fewer. Firms that remain listed throughout the entire sample period are therefore more resilient than those that enter and exit the market.

2.3 The profile of distressed firms

Table 3 presents descriptive statistics for some indicators of each category, along with mean values for distressed and non-distressed firms and the univariate AUROC of each

indicator. Several differences emerge between the two groups. Firms that enter financial distress show lower profitability, lower liquidity, and a larger short-term debt than their non-distressed counterparts. EBIT over total assets averages 10.4% among healthy observations and −1.0% among distressed ones, and the operating margin averages 21.4% against −3.2%. Distressed firms also display lower liquidity: the cash ratio falls from 0.49 to 0.37 and working capital over assets from 11.5% to 4.2%, indicating a weaker capacity to absorb financial shocks. Coverage indicators show the expected pattern: distressed firms face a higher cost of debt (7.5% against 4.9%), consistent with a reduced ability to service financial obligations out of operating earnings. For activity ratios, both asset turnover and fixed-asset turnover indicate lower operating efficiency among distressed firms. Finally, in the growth and size category, distressed firms tend to be smaller (12.2 against 13.0 in log assets) and younger (8.6 against 10.6 years).

The AUROC results reveal large differences in the discriminatory power of financial ratios. Profitability measures exhibit the highest discriminatory power , with univariate AUROC values of 0.93 for EBIT over assets and 0.88 for the operating margin. In contraste, leverage indicator reaches only 0.53, while liquidity, activity and size indicators lie in between, between 0.59 and 0.75.

Table 3: Descriptive statistics of the leading indicators by category (winsorised at 5% per tail)

						Mean by status		
Category	Indicator	N	Mean	SD	Median	Healthy	Distress	AUROC
Profitability	EBIT to assets (X_3)	3,103	0.082	0.089	0.067	0.104	-0.010	0.932
	Operating margin	2,925	0.175	0.204	0.134	0.214	-0.032	0.876
Liquidity	Cash ratio	3,099	0.466	0.706	0.141	0.488	0.374	0.657
	Working capital to assets (X_1)	3,103	0.101	0.172	0.074	0.115	0.042	0.628
Leverage	Debt maturity structure	3,099	0.566	0.281	0.559	0.562	0.586	0.527
Coverage	Cost of debt	3,098	0.054	0.049	0.039	0.049	0.075	0.639
Activity	Asset turnover (X_5)	3,103	0.580	0.462	0.469	0.644	0.304	0.748
	Fixed-asset turnover	3,035	1.698	2.123	0.821	1.821	1.110	0.674
Growth/Size	Size (log assets)	3,104	12.832	1.581	12.844	12.973	12.227	0.626
	Firm age	3,104	10.228	6.148	10.000	10.607	8.598	0.599

Notes: estimation sample (permanence ≥ 5 years, 2000–2025), 3,195 firm-years. Source: authors' calculations.

3 Empirical Strategy

3.1 Estimation and testing sample

All models are estimated over the 2000–2019 period. The out of sample testing covers 2020-2025, which includes the COVID-19 shock, providing a test for model performance under stressed conditions. Using the same evaluation period across all models ensures

comparability of their predictive performance.

To assess performance, we report the area under the ROC curve (AUROC), overall accuracy and the two classification errors, all computed on the test set. *Type I error* is a healthy firm classified as distressed, a false alarm, and a *Type II error* is a distressed firm classified as healthy, a missed event. The distinction matters for how these models would be used. From a supervisory perspective, the Type II error is more concern, since it corresponds to a firm heading into distress that the indicator fails to flag.

3.2 Candidate indicators and variable selection

Financial distress is not observable before default or formal insolvency, so the literature reads a firm's condition from its financial statements. This works because distress is a gradual process: as a firm deteriorates, the symptoms show up in its performance before they become visible as a legal or market outcome. Following Beaver (1966), Pinches et al. (1973) and Chen and Shimerda (1981), we group the candidate ratios into six economic dimensions: profitability, liquidity, leverage, coverage, activity, and growth and size.

I. Profitability. Profitability measures resource generation: how much cash flow a firm can create to service debt and replenish equity. Given that continuous losses are one of the main signals of financial distress, the literature usually treats profitability as the leading indicator of a firm's capacity to remain viable. The expected association with financial distress is negative: more profitable firms should be less likely to experience distress.

II. Liquidity. Liquidity determines whether temporary cash-flow shortfalls translate into distress. These ratios capture the burden of short-term obligations relative to liquid resources. The expected relationship with financial distress is also negative: firms with higher liquidity ratios should be less exposed to distress.

III. Leverage. Leverage is the channel through which a shock becomes a threat to solvency. Because debt creates fixed payment obligations, highly leveraged firms require larger and more stable cash flows to remain viable. A decline in profit is more likely to result in financial distress for a highly leveraged firm than for a firm with a lower debt burden. The expected association between leverage and distress is positive.

IV. Coverage. Coverage ratios connect operating earnings directly to debt service. They capture a firm's capacity to meet its financial obligations out of internally generated resources. Low coverage ratios indicate that operating earnings provide only limited protection against interest expenses, increasing the likelihood that adverse shocks will compromise the firm's ability to service its debt. Stronger coverage is expected to be associated with lower distress risk.

V. Activity. Activity ratios measure how efficiently firms use their assets to generate revenue. Lower asset utilization is linked to operational inefficiencies and weak cash generation, which can reduce a firm's ability to meet its financial obligations and increase the likelihood of distress. Higher activity is generally associated with lower distress risk.

VI. Growth and size. This dimension captures firm characteristics that represent resilience to adverse shocks. Firm size may reflect greater diversification and better access to external financing. Firm age captures organisational maturity and experience, and also identifies the firm's position along its life-cycle. Unlike the previous dimensions, the expected relationship between growth and financial distress is ambiguous. Low or declining growth may signal weakening demand and deteriorating operating conditions, increasing vulnerability to distress. In contrast, rapid growth may reflect overexpansion and create financial pressures that increase future vulnerability.

The six dimensions capture complementary aspects of one process and are not independent causes of distress.

Variable selection and dimensionality reduction. The set of predictors assembled from the available data contains 39 indicators across the six categories, listed in the Appendix. Two of them are excluded , before any selection takes place: the interest coverage ratio and financial debt over EBIT. Both are functions of the same earnings and interest figures that define the event of distress in equation (1), so including them would mechanically predict the outcome. Of the remaining indicators, those defined for at least 90% of the training window are retained, leaving **26 valid candidates**. Following Hernandez Tinoco and Wilson (2013), candidates are winsorised at the 5th and 95th percentiles of the training window to limit the effect of outliers, and each predictor enters lagged one year.

In order to reduce the number of predictors, limit multicollinearity, and avoid redundancy within categories, we consider two alternative procedures. The first set is chosen by the data. We apply the LASSO (Tibshirani, 1996) to every candidate predictor and let the penalty determine which variables are retained, solving the ℓ_1-penalised logit

$$\hat{\boldsymbol{\beta}}(\lambda) \;=\; \arg\min_{\boldsymbol{\beta}} \Big\{ -\tfrac{1}{n} \textstyle\sum_{i,t} \ell_{it}(\boldsymbol{\beta}) \;+\; \lambda \|\boldsymbol{\beta}\|_1 \Big\}, \tag{3}$$

where $\ell_{it}(\boldsymbol{\beta})$ is the Bernoulli log-likelihood of the distress outcome for firm-year it, $\|\boldsymbol{\beta}\|_1 = \sum_k |\beta_k|$ is the ℓ_1 norm, n is the number of training firm-years, and $\lambda \geq 0$ is the penalty strength, chosen by 5-fold stratified cross-validation on the training window under the minimum-loss rule. This retains **19 variables** and defines the *LASSO logit* model.

The second approach is a category-structured alternative, implemented as a three-stage reduction of the 26 valid candidates.

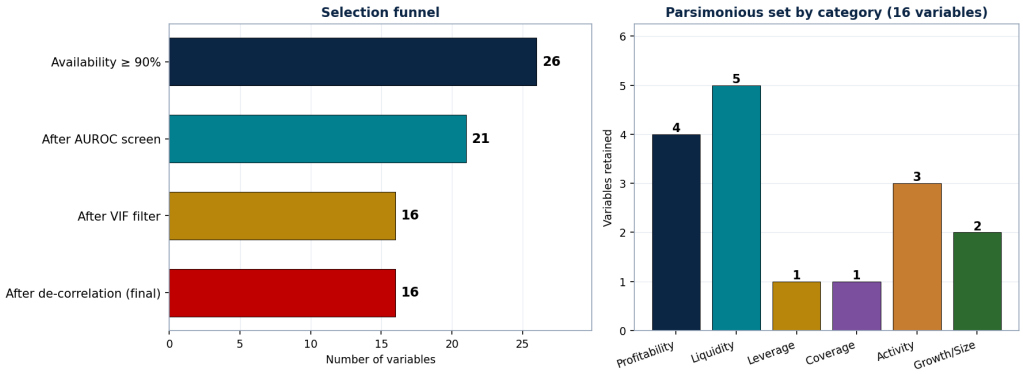
Stage 1 — univariate screening. Each lagged predictor is tested for discriminatory power by its univariate AUROC on the training window, with a 500-replication bootstrap interval, retaining those with AUROC ≥ 0.55. Twenty one indicators clear the threshold.

Stage 2 — intra-category redundancy. Within each category, survivors are ordered by their AUROC and those with high multicollinearity are excluded: any indicator with a per-category variance-inflation factor (VIF) above 5 is dropped in favour of the better-discriminating same-category alternative. This stage removes five indicators and leaves **16**.

Stage 3 — global de-correlation. Across the retained set, any pair with absolute Pearson correlation above 0.85 is resolved by dropping the lower-AUROC member. No pair exceeds the threshold at this point, so the set is unchanged. The result is the *parsimonious set* of **16 variables**, spanning all six categories.

Figure 1 summarises the funnel:

Figure 1: Multistage variable-selection procedure and composition of the predictor set



3.3 Distress models

We estimate three groups of models: The Linear Discriminant Analysis (LDA) , a set of logit specifications with alternative predictor sets and two dynamic specifications. Table 4 summarises the nine models, their predictors and their classification rules.

3.4 LDA and recalibration of the Altman model

The initial model is the linear discriminant analysis (LDA), which assigns each firm a score based on

$$Z_i \;=\; \hat{\boldsymbol{w}}'\boldsymbol{x}_i, \qquad \hat{\boldsymbol{w}} \;=\; \hat{\boldsymbol{\Sigma}}^{-1}(\bar{\boldsymbol{x}}_H - \bar{\boldsymbol{x}}_D)\,, \tag{4}$$

where $\boldsymbol{x}_i$ is the vector of predictors—in our case the five Altman ratios: working capital, retained earnings and EBIT (each over total assets), the market value of equity over total liabilities, and asset turnover (X_1–X_5)—for firm i; $\bar{\boldsymbol{x}}_H$ and $\bar{\boldsymbol{x}}_D$ are the ratio

Table 4: Summaries the models

Model	Predictors	Estimator / cutoff	Regime
Altman 1968-benchmark	X_1–X_5,	LDA; $Z < 1.81$	A
LDA (matched)	X_1–X_5	LDA; centroid midpoint	A
LDA (pooled)	X_1–X_5	LDA; centroid midpoint	B
Logit base	X_1–X_5	MLE logit; Youden	B
Logit all	all indicators	MLE logit; Youden	B
Logit LASSO	19 LASSO-selected ^{s}	MLE logit; Youden	B
Logit parsimonious	16 reduction-set ^{s}	MLE logit; Youden	B
Panel logit + macro	LASSO set + 6 macro covariates	Pooled logit, firm-clustered SE; Youden	B
Hazard + macro	LASSO set + macro + $\ln(1+\text{age})$	Discrete-time hazard (Shumway, 2001); Youden	B

Regimes: A = matched cross-section (one size-matched observation per firm); B = full pooled sample with firm-clustered standard errors.

means of the healthy and distressed groups; $\hat{\boldsymbol{\Sigma}}$ is the pooled within-group covariance matrix; and $\hat{\boldsymbol{w}}$ is the Fisher direction that maximises between- to within-group variance under equal priors. The score is oriented so that a higher Z_i indicates a healthier firm; a firm is classified as distressed when $Z_i < c$.

Under equal priors, we set the cutoff at the midpoint of the group centroids,

$$c^*_{\text{LDA}} \;=\; \tfrac{1}{2}\big(\bar{Z}_H + \bar{Z}_D\big)\,, \tag{5}$$

where $\bar{Z}_H$ and $\bar{Z}_D$ are the mean discriminant scores of the healthy and distressed groups on the training sample. The classical benchmark for the LDA model is the Altman (1968) Z-score, with its original coefficients and the $Z < 1.81$ cutoff.

The recalibration of the LDA model is performed under two sample designs. *Regime A (matched cross-section)* keeps one observation per firm—the last training firm-year with a defined outcome—with each distressed firm matched one-to-one to its nearest healthy neighbour in firm size, approximating the original design of Altman (1968). *Regime B (full pooled sample)* uses all firm-years with standard errors clustered at the firm level (Shumway, 2001; Petersen, 2009).

3.5 Logit specifications

We also consider logit models, given their proven accuracy and predictive power for financial-distress forecasting (Ohlson, 1980; Campbell et al., 2008; Hernandez Tinoco and Wilson, 2013). We estimate the following specification:

$$\mathbb{P}(\text{distress}_{it} = 1 \mid \boldsymbol{x}_{i,t-1}) \;=\; \Lambda\big(\alpha + \boldsymbol{\beta}'\boldsymbol{x}_{i,t-1}\big), \qquad \Lambda(u) = \frac{e^u}{1+e^u}, \tag{6}$$

where distress_{it} is the binary outcome for firm i in year t; $\boldsymbol{x}_{i,t-1}$ is the vector of one-year-lagged predictors; α is the intercept and $\boldsymbol{\beta}$ the coefficient vector; and $\Lambda(\cdot)$ is the logistic function. The model is estimated by maximum likelihood on the full pooled training sample with standard errors clustered at the firm level.

Given the large number of candidate predictors, we estimate it on four predictor sets: the five original Altman ratios (the *Altman-ratio logit*), which isolates the gain from changing estimator while holding the information fixed; every candidate that meets the availability floor (the *full logit*); the 19 variables chosen by the LASSO; and the 16 variables of the parsimonious set.

3.6 Dynamic specifications

We also consider two additional specifications, a panel logit and a hazard model, to test whether dynamic structure and macroeconomic variables improve the forecasting power of the models. Following Hernandez Tinoco and Wilson (2013) and Nam et al. (2008), we add six macroeconomic variables (real GDP growth, the real interest rate, the credit spread, inflation, FX depreciation, and the terms of trade), together with the LASSO predictor set, all lagged one year. The *panel logit* is Equation (6) estimated over the firm-year panel with firm-clustered standard errors. The *discrete-time hazard* model adds a baseline term $\gamma \ln(1 + \text{age}_{i,t-1})$ inside the index of Equation (6), where $\text{age}_{i,t-1}$ is the lagged firm age (years since the firm first appears in the panel) and γ is the baseline-hazard coefficient (Shumway, 2001; Beck et al., 1998).

Classification and threshold. Each logit model requires a threshold to convert its predicted probability into a binary classification. For each model, we report the cutoff that maximizes the Youden index:

$$J(c) = \text{Sens}(c) + \text{Spec}(c) - 1, \qquad c^* = \arg\max_c J(c), \tag{7}$$

where $\text{Sens}(c)$ is the sensitivity, or true positive rate, defined as the share of distressed firms correctly classified as distressed at threshold c, and $\text{Spec}(c)$ is the specificity, or true negative rate, defined as the share of non-distressed firms correctly classified as non-distressed . The index ranges from 0, for a classifier no better than random, to 1, for perfect classification. We use the Youden index because it gives equal weight to both types of classification error, missed distress cases and false alarms, and does not depend on the sample share of distressed firms. The cutoff is set on the training window and then applied to the test period.

3.7 The aggregate stress index and its decomposition

Following the aggregation index literature (Galán, 2021), the firm-level probabilities are aggregated into a corporate financial stress indicator. We report an *extensive* margin (the share of firms in stress) and an *intensive* margin (an asset- or debt-weighted average of the probabilities),

$$\mathrm{FSI}_t^{ext} = \frac{1}{N_t}\sum_i \mathbf{1}\{\hat{p}_{it} > c^*\}, \qquad \mathrm{FSI}_t^{int} = \sum_i w_{it}\,\hat{p}_{it}, \quad w_{it} = \frac{A_{it}}{\sum_j A_{jt}}, \tag{8}$$

where $\hat{p}_{it}$ is the fitted distress probability of firm i in year t; c^* is the training-frozen cutoff; N_t is the number of firms in year t; $\mathbf{1}\{\cdot\}$ is the indicator function; A_{it} is total assets; and w_{it} is the asset weight (the debt-weighted variant replaces A_{it} by total liabilities).

Category decomposition. To identify the financial dimensions underlying the fitted distress probabilities, we apply a grouped Shapley decomposition (Shapley, 1953; Owen, 1977). The explanatory variables are grouped into the six categories described before, and the contribution of each category is computed as its average marginal effect on the predicted probability across all possible combinations of categories. The resulting decomposition is additive, so that the fitted probability can be expressed as the sum of a baseline probability and the contributions of the six financial dimensions.

Let $g_c(\boldsymbol{x}_i) = \sum_{k\in c}\hat{\beta}_k\, x_{k,it}$ be the parsimonious-logit index contribution of category c, and for a coalition $S \subseteq \mathcal{C}$ of the $m = 6$ categories define

$$p_i(S) \;=\; \mathbb{E}_b\,\Lambda\Big(\hat{\alpha} + \sum_{c\in S} g_c(\boldsymbol{x}_i) + \sum_{c\notin S} g_c(\boldsymbol{x}_b)\Big), \tag{9}$$

where Λ is the logistic link, the categories in S take firm i's values, and the remaining categories are drawn from a background sample b (the training firm-years); the expectation $\mathbb{E}_b$ averages over that background. The Shapley contribution of category c to firm i's probability is

$$\phi_{ic} \;=\; \sum_{S\subseteq\mathcal{C}\setminus\{c\}} \frac{|S|!\,(m-|S|-1)!}{m!}\,\big[p_i(S\cup\{c\}) - p_i(S)\big]. \tag{10}$$

These contributions are measured in *probability units* and satisfy the efficiency (local-accuracy) property

$$p_i \;=\; \bar{p} + \sum_{c\in\mathcal{C}}\phi_{ic}, \qquad \bar{p} \;=\; \mathbb{E}_b\,\Lambda\big(\hat{\alpha} + \textstyle\sum_c g_c(\boldsymbol{x}_b)\big), \tag{11}$$

so each firm's fitted distress probability decomposes *exactly* into a common baseline $\bar{p}$—the average predicted probability, close to the sample distress rate—plus a signed contribution

per category: a positive ϕ_{ic} pushes the firm toward distress, a negative one protects it.

The decomposition aggregates linearly. With asset weights $w_{it} = A_{it}/\sum_j A_{jt}$, the systemic index inherits the same split,

$$\mathrm{FSI}_t^{int} \;=\; \bar{p} + \sum_{c \in \mathcal{C}} C_{ct}, \qquad C_{ct} \;=\; \sum_i w_{it}\, \phi_{ic}, \tag{12}$$

so C_{ct} measures, in probability points, how much category c adds to aggregate corporate stress at date t.

4 Results

4.1 Altman model and the LDA recalibration

Table 5 reports the results of the Altman model recalibration under the two estimation regimes. In addition to the estimated discriminant coefficients, the table presents the endogenous classification cutoffs and the corresponding in-sample performance metrics.

Table 5: Recalibration of the Altman discriminant: coefficients, cutoff and in-sample fit

	Original (1968) fixed coeff.	Recalibrated matched (A)	Recalibrated pooled (B)
X_1 WC/assets	1.200	2.347	-0.032
X_2 RE/assets	1.400	4.413	3.472
X_3 EBIT/assets	3.300	26.399	14.359
X_4 Equity/liab.	0.600	-0.134	-0.080
X_5 Asset turnover	1.000	-0.021	0.585
Cutoff	1.81	1.235	0.848
In-sample AUROC	0.666	0.940	0.911
In-sample Type I error	0.312	0.163	0.200
In-sample Type II error	0.370	0.082	0.122
In-sample overall accuracy	0.677	0.878	0.814
N (training)	2,495	98	2,495

Notes: the original model keeps the 1968 coefficients and the $Z < 1.81$ rule. In-sample metrics are computed on the training window (2000–2019).Source: authors' calculations.

The first result is related to the estimated coefficient. The profitability ratios (X_2 and X_3) retain the positive signs found in the original Altman model and receive larger weights after recalibration, suggesting that accumulated earnings and operating profitability are the strongest discriminators between distressed and non-distressed firms in the Peruvian sample. Working capital over assets (X_1) also remains positive in the matched specification, , although its contribution becomes negative in the pooled estimation.

By contrast, the leverage ratio (X_4), switches sign in both recalibrated models and the The activity ratio (X_5) displays unstable signs across estimation regimes. Those results may reflect the high correlation among the predictors .

The in-sample classification results point to a large improvement relative to the Altman specification: the AUROC rises from 0.666 under the original coefficients to 0.940 in the matched specification and 0.911 in the pooled one, with Type I errors falling from 31.2% to 16.3% and 20.0% and Type II errors from 37.0% to 8.2% and 12.2%.

For the out of sample results, Table 6 presents the performance of the 3 models :

The results confirm the benefits of recalibrating the model with Peruvian data. While the original Altman model achieves an AUROC of 0.779, both recalibrated specifications reach higher values: 0.949 for the matched specification and 0.936 for the pooled one.

Table 6: Out-of-sample performance: Original model versus recalibrated

Model	AUROC	Type I	Type II	Accuracy	Cutoff
Altman (1968), original	0.779	0.410	0.204	0.625	1.81
Recalibrated (matched)	0.949	0.159	0.061	0.857	1.235
Recalibrated (pooled)	0.936	0.143	0.143	0.857	0.848

Notes: common test set (2020–2025), 581 firm-years of which 98 are distress events; each model at its own cutoff (original at $Z < 1.81$, recalibrated at the centroid midpoint frozen on the training window). Source: authors' calculations.

In terms of classification errors, the matched recalibration reduces the Type I error from 41.0% to 15.9% and the Type II error from 20.4% to 6.1%, while the pooled specification achieves an even lower Type I error (14.3%) at the cost of a higher Type II error (14.3%). Both recalibrations reach a similar overall accuracy of 85.7%, against 62.5% for the original model.

Overall, these results suggest that the recalibrated models capture features of financial distress that are specific to the Peruvian corporate sector, improving classification accuracy and reducing both types of error. Between the two specifications, the matched model achieves the highest discriminatory power and lower type II error.

4.2 Logit specifications

The next step is to change the model, from LDA to Logit, to measure if the forecasting power improves when changing specification, using the same five financial ratios.

Table 7: The Altman-ratio logit

Variable	Coef.	Std. err.	z	Odds ratio
Constant	0.398**	0.198	2.01	1.489
X_1	-0.240	0.738	-0.33	0.787
X_2	-1.651***	0.466	-3.54	0.192
X_3	-39.877***	2.804	-14.22	0.000
X_4	-0.011	0.027	-0.42	0.989
X_5	-0.958***	0.312	-3.07	0.384

Notes: dependent variable is the distress indicator of Equation (1); all predictors enter lagged one year. Standard errors are clustered at the firm level. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. $N = 2{,}495$. Source: authors' calculations.

The estimated coefficients reported in Table 7 have the expected negative sign for all five ratios, indicating that higher values are associated with a lower probability of being classified as distressed, although they differ in magnitude and statistical significance. X_2 (retained earnings over assets), X_3 (EBIT over assets) and X_5 (asset turnover) are significant at the 1% level, suggesting that these dimensions contain relevant information for distinguishing distressed from non-distressed firms in a multivariate specification.

Among them, X_3 exhibits the largest coefficient in absolute value, suggesting that operating profitability is the most important predictor of financial distress among the five Altman ratios.

The out-of-sample results in Table 8 show that using the same five Altman ratios in a logistic specification produces slightly better performance than the recalibrated LDA models. The logit model achieves an AUROC of 0.953, above both LDA models. It also achieves the highest overall accuracy, 86.6%, while maintaining relatively low classification errors: its Type I error of 14.7% is below that of the matched LDA (15.9%), although its Type II error of 7.1% is above. Overall, the gains from changing the estimator are modest.

Table 8: Discriminant analysis versus logit: out-of-sample discriminatory power

Model	AUROC	Type I	Type II	Accuracy	Cutoff
Recalibrated LDA (matched)	0.949	0.159	0.061	0.857	1.235
Recalibrated LDA (pooled)	0.936	0.143	0.143	0.857	0.848
Altman-ratio logit	0.953	0.147	0.071	0.866	0.216

Notes: common test set (2020–2025), 581 firm-years; Youden cutoff for the logit, centroid midpoint for the two LDA specifications. Source: authors' calculations.

4.3 Expanding the set of predictors

The next step is to increase the set of predictors in order to assess the potential gains in accuracy and discriminatory power. We estimate a single logit model that includes all valid predictors (the 26 indicators from the six economic categories), which we refer to as the *full logit*. Additionally, we estimate a LASSO logit, which lets the data select the predictors: it retains 19 of the 26 candidates (six profitability, three liquidity, three leverage, one coverage, four activity, and two growth and size). Table 15 in the Appendix reports the regression estimates for both models, together with the variables each one selects.

Table 9 shows the out of sample performance of these 2 models, compare to the original logit model with the 5 Altman ratios.

Table 9: Expanding the predictor set: out-of-sample performance

Model	AUROC	Type I	Type II	Accuracy	Cutoff
Altman-ratio logit	0.953	0.147	0.071	0.866	0.216
Full logit	0.967	0.116	0.041	0.897	0.144
LASSO logit	0.966	0.112	0.041	0.900	0.147

Notes: common test set (2020–2025), 581 firm-years of which 98 are distress events; Youden cutoff frozen on the training window. The full logit uses 26 predictors and the LASSO logit 19. Source: authors' calculations.

The results in Table 9 show that expanding the predictor set improves out-of-sample performance relative to the baseline logit . The full specification raises the AUROC from 0.953 to 0.967 and overall accuracy from 86.6% to 89.7%. The improvement is reflected in both classification errors: the Type I error falls from 14.7% to 11.6% and the Type II error from 7.1% to 4.1%. These results indicate that the additional financial ratios contain information about financial distress that is not captured by the five ratios of the Altman specification.

The LASSO specification delivers similar performance with seven fewer variables. Its AUROC of 0.966 is close to the full model, while its overall accuracy is slightly higher, at 90.0%, and its Type I error slightly lower, at 11.2%. The two models attain the same Type II error of 4.1%.

Overall, the results suggest that most of the gains from expanding the predictor set can be achieved without retaining all available variables, and that the dimensionality of the predictor set can be reduced while retaining the model's predictive power.

4.4 The parsimonious specification

Using the reduced set of variables from the dimensionality reduction procedure, we estimate a new logit model ("Parsimonious logit") on the same training sample (2000–2019) and evaluate it on the common test sample (2020–2025).

Table 10: Out-of-sample discriminatory power of alternative logit specifications

Model	k	AUROC	Type I	Type II	Accuracy	Cutoff
Altman-ratio logit	5	0.953	0.147	0.071	0.866	0.216
LASSO logit	19	0.966	0.112	0.041	0.900	0.147
Full logit	26	0.967	0.116	0.041	0.897	0.144
Parsimonious logit	16	**0.969**	**0.101**	**0.020**	**0.912**	0.139

Notes: Cutoffs are selected according to the Youden criterion on the training window. Source: authors' calculations.

Table 10 shows that the parsimonious model matches or exceeds every alternative model on the selected metrics while using 16 predictors. Its AUROC reaches 0.969, against 0.967 for the full logit and 0.966 for the LASSO, and its overall accuracy is 91.2%, against 89.7% and 90.0%. The parsimonious specification achieves a Type I error of 10.1%, compared with 11.6% for the full logit and 11.2% for the LASSO. Its Type II error is 2.0%, half the 4.1% obtained by both of the larger specifications: of the 98 distress events in the test sample, the parsimonious model misses two, while the LASSO and full logit each miss four. These results suggest that the selection procedure preserve discriminatory whith fewer predictors, by removing redundant variables and reducing the multicollinearity that distorts the weights of the specification.

4.5 Dynamic models: panel logit and hazard

The next exercise extends the analysis by incorporating the time dimension through two dynamic specifications: a panel logit and a discrete-time hazard model. In addition to firm-level financial indicators, both specifications include a set of macroeconomic variables. The objective is to assess whether accounting for the temporal dimension and macroeconomic conditions improves discriminatory power.

Table 11: Static versus dynamic models: out-of-sample performance

Model	AUROC	Type I	Type II	Accuracy	Cutoff
Parsimonious logit	0.969	0.101	0.020	0.912	0.139
LASSO logit	0.966	0.112	0.041	0.900	0.147
Panel logit + macro	0.965	0.135	0.010	0.886	0.147
Hazard + macro	0.963	0.116	0.010	0.902	0.166

Notes: common test set (2020–2025). Both dynamic specifications use the LASSO predictor set augmented with six lagged macroeconomic variables; the hazard model adds the baseline term $\ln(1 + \text{age})$. Source: authors' calculations.

Table 11 compares the out-of-sample performance of the dynamic specifications with that of the best-performing static models. The first result is that incorporating the time dimension and macroeconomic variables does not improve overall discriminatory power. The parsimonious logit remains the best-performing specification, with an AUROC of 0.969 and an overall accuracy of 91.2%, against 0.965 and 88.6% for the panel logit and 0.963 and 90.2% for the hazard model. The dynamic models do, however, improve the detection of distress events. The panel logit and the hazard model each reduce the Type II error from 2.0% in the parsimonious specification to 1.0%: they miss one distress event out of 98, against two for the parsimonious logit and four for the LASSO logit. This gain comes at the cost of a higher Type I error. The panel logit raises it from 10.1% to 13.5%, and the hazard model to 11.6%. The reduction in missed distress events is thus offset by a larger number of false alarms, and overall accuracy declines in both cases.

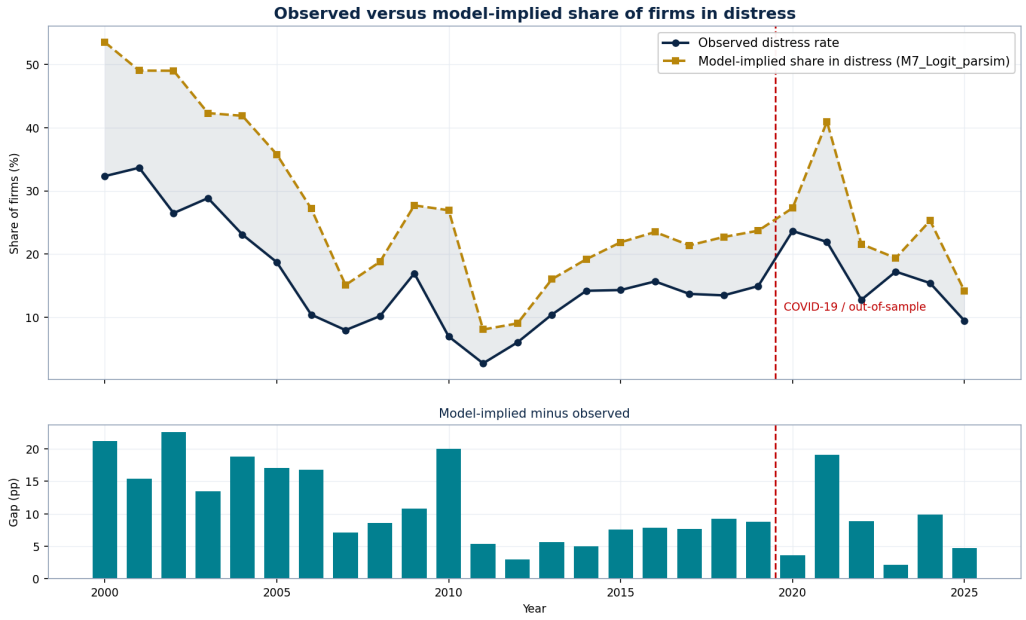
The results suggest that the information contained in the macroeconomic variables and in the temporal structure of the panel adds limited predictive content beyond that already captured by firm-level financial indicators, although it slightly improves the detection of distress events at the cost of more false alarms.

4.6 The aggregate stress indicator

Figure 2 compares the average predicted probability of financial distress from the parsimonious model with the observed distress rate over the period 2000–2025.

The two series display similar dynamics throughout the sample and identify the same episodes of corporate vulnerability. Both measures indicate elevated distress during the

Figure 2: Estimated probability and observed financial distress rate



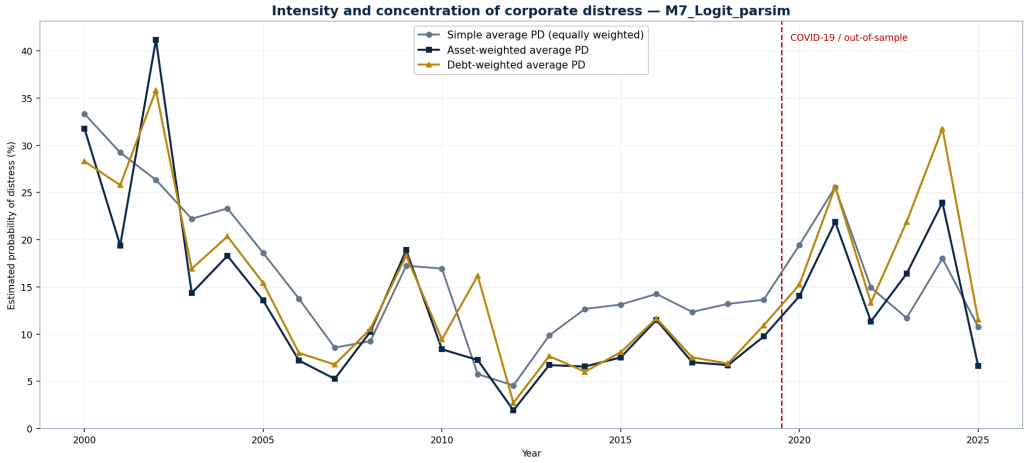
early 2000s, an improvement during the commodity-boom years, and a deterioration following the COVID-19 shock. Over the full period the observed distress rate averages 16.2%, with a minimum of 2.7% in 2011 and a maximum of 33.7% in 2001; while the model-implied share averages 27.0%, ranging from 8.0% in 2011 to 53.5% in 2000.

The predicted share of firms classified as being at risk is consistently higher than the observed distress rate, with an average difference of 10.8 percentage points over the sample. This level difference is a consequence of the classification rule, the cutoff is selected according to the Youden criterion, which gives equal weight to false-positive and false-negative classifications and therefore does not seek to reproduce the observed incidence of distress. In this sense, the model-implied share should be interpreted as an alert rate: the proportion of firms whose financial characteristics place them in the region associated with elevated distress risk, rather than the proportion of firms that will necessarily realize the distress outcome.

Building on this evidence, we aggregate the firm-level predicted probabilities into a corporate financial stress indicator using alternative weighting schemes that account for differences in firms' economic size and debt exposure. Following Galán (2021), we construct two complementary intensive margins based on asset- and debt-weighted averages of predicted distress probabilities. The asset-weighted indicator captures the prevalence of predicted distress relative to the size of firms, whereas the debt-weighted indicator provides a measure of potential creditor exposure.

Figure 3 shows the evolution of the two intensive margins. Both display a similar cyclical pattern, with elevated stress in the early 2000s, a sustained decline during the commodity-boom years, and a renewed increase from the mid-2010s onward. Both margins

Figure 3: Intensive margins of the corporate financial stress indicator



reach their maximum in 2002 and their minimum in 2012. The increase around 2008 coincides with the global financial crisis, while the subsequent deterioration is consistent with tighter global financial conditions and the decline in commodity prices observed in Peru from 2013 onward. The deterioration became pronounced during the COVID-19 pandemic. The asset-weighted indicator rose from 9.7% in 2019 to 21.9% in 2021, while the debt-weighted measure increased from 10.9% to 25.6%. This episode coincided with the severe contraction in economic activity in 2020 and the subsequent recovery. Following a temporary decline in 2022, stress increased again in 2023 and 2024, in line with the normalisation of monetary policy and the higher global interest-rate environment. In 2024 the debt-weighted indicator reached 32.0%, compared with 24.1% for the asset-weighted measure (Table 13), indicating that predicted distress was relatively more concentrated among firms with greater debt exposure. By 2025 both measures had declined sharply, to 6.6% and 11.5% respectively. Overall, the intensive margins capture the prevalence of predicted distress while accounting for its concentration among firms with greater economic size and liability exposure.

The aggregate indicators presented so far show when corporate financial stress increased and how its intensity varied according to firms' economic size and liability exposure. However, they don't indicate which financial dimensions are most closely associated with those changes. To answer this, we decompose the predicted distress probabilities into the six financial categories—profitability, liquidity, leverage, coverage, activity, and growth and size, using the grouped Shapley decomposition.

The decomposition is considered from two perspectives. First, the cross-sectional results identify the financial dimensions that distinguish firms with high predicted distress from those with lower predicted distress. Second, the time-series decomposition shows how the contribution of these dimensions changes across episodes of corporate stress.

Table 12 reports the cross-sectional decomposition. Relative to the baseline probability of 17.13%, profitability accounts for the largest share of the difference in predicted distress

between the two groups. It contributes an average of 36.46 percentage points in distressed firm-years and −7.42 percentage points in healthy ones, a gap of 43.88 percentage points that represents 80.8% of the total. Activity follows with a gap of 3.67 percentage points, or 6.8%, and coverage with 3.29 percentage points, or 6.1%.Liquidity accounts for 4.8%, while growth/size and leverage together explain less than 2%. Taken together, profitability and activity account for approximately 87% of the difference in predicted distress between healthy and distressed firms.

Table 12: Category importance in the fitted distress probability-Shapley values

	Mean $\|\phi\|$	Mean contribution			Share of
Category	(pp)	Healthy	Distress	Gap (pp)	gap (%)
Profitability	17.44	-7.42	36.46	43.88	80.8
Activity	2.76	-0.36	3.31	3.67	6.8
Coverage	4.13	-1.06	2.23	3.29	6.1
Liquidity	2.57	-0.31	2.27	2.58	4.8
Growth/Size	2.02	-0.60	-0.00	0.60	1.1
Leverage	1.28	0.11	0.41	0.31	0.6

Notes: grouped Shapley values of the parsimonious logit, in probability points, relative to a baseline probability of $\bar{p} = 17.13\%$. "Mean $|\phi|$" is the average absolute contribution across firm-years; the remaining columns compare the average contribution in healthy and in distressed firm-years and the share of that gap explained by each category. Source: authors' calculations.

Table 13 decomposes the aggregate stress indicator over time, reporting the average contribution of each category by sub-period and for a set of relevant years.

Table 13: Category contributions to the aggregate stress index over time

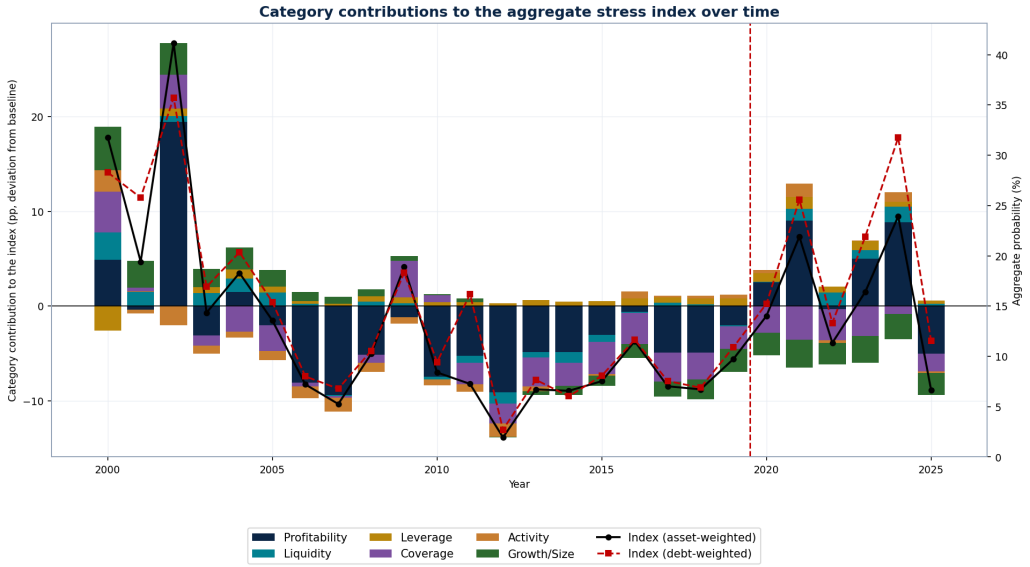
	Category contribution (pp)						Index (%)	
	Profit.	Cover.	Activ.	Liquid.	Growth	Lever.	Assets	Debt
Period averages								
2000–2007	-0.41	+0.09	-0.37	+0.82	+2.25	+0.27	19.8	20.5
2008–2011	-6.37	+0.33	-0.63	-0.11	+0.36	+0.54	11.2	13.7
2012–2019	-5.87	-2.94	+0.06	-0.43	-1.37	+0.62	7.2	7.7
2020–2025	+1.83	-2.68	+0.53	+0.91	-2.72	+0.79	15.8	20.0
Selected years								
2002 (sample peak)	+17.79	+3.47	-1.76	+0.59	+3.20	+0.84	41.3	36.2
2009 (global crisis)	-2.74	+3.84	-0.51	+0.25	+0.41	+0.67	19.0	18.4
2012 (sample trough)	-10.63	-2.19	-1.35	-1.24	-0.09	+0.31	1.9	2.7
2020 (COVID-19)	+0.98	-2.89	+0.47	+0.07	-2.58	+0.91	14.1	15.3
2024	+7.42	-0.89	+1.14	+1.60	-2.83	+0.53	24.1	32.0

Notes: asset-weighted Shapley contribution of each category, in probability points relative to the baseline probability of $\bar{p} = 17.13\%$; the six contributions plus the baseline add up exactly to the asset-weighted index in the penultimate column, and the last column reports the same index weighted by liabilities. The index is recomputed on the sample used by the decomposition, which from 2008 onward coincides with the one used for the annual series of Section 4.6 to within 0.2 percentage points. Source: authors' calculations.

During 2000–2007 the aggregate index remained high, with positive contributions from growth and size and only limited contributions from profitability and coverage. The decline in stress between 2008 and 2011 was driven primarily by an improvement in profitability, whose contribution fell to −6.37 percentage points. This pattern became more pronounced during 2012–2019, when profitability and coverage were both negative, reducing the index to an average of only 7.2%. In contrast, the post-pandemic period is characterised by a reversal of profitability, which becomes a positive contributor to aggregate stress.

Figure 4 presents the annual decomposition of the aggregate indicator. It highlights the dominant role of profitability throughout the sample: fluctuations in its contribution closely track the major changes in aggregate corporate stress, while the remaining categories generally make smaller contributions.

Figure 4: Annual decomposition of the aggregate stress indicator by category



Overall, The Shapley decomposition indicates that profitability is the main financial dimension underlying corporate stress in Peru. Both the cross-sectional and the time-series results show that profitability accounts for most of the variation in predicted distress, while the remaining categories play more limited and episode-specific roles.

5 Conclusion

This paper develops a financial stress indicator for Peruvian non-financial corporations using firm-level financial statements over the period 2000–2025. The analysis combines distress-prediction models, variable-selection techniques, and an aggregate stress-index framework in order to construct a measure of corporate vulnerability.

Several conclusions emerge. First, the direct application of the original Altman (1968)

model to Peruvian firms leads to large classification errors, with an out-of-sample AUROC of 0.78 and a Type I error above 40%. Recalibrating the model with local data raises the AUROC to 0.94–0.95 and, under the matched design, cuts both errors by more than half, marking the importance of adapting financial-distress models to specific countries and firms.

Second, expanding the information set beyond the original Altman ratios produces additional gains. The multi-stage selection procedure identifies a parsimonious set of 16 indicators. The resulting logit specification attains an out-of-sample AUROC of 0.97 and a Type II error of 2.0%, outperforming both the all variable specification and the LASSO selection while using fewer predictors . While panel-logit and hazard models reduce the Type II error again, to 1.0%, they do it at the cost of more false alarms, and do not improve overall discriminatory power.

Third, firm-level distress probabilities can be aggregated into a measure of corporate-sector vulnerability. The resulting indicator captures the major episodes of stress observed over the sample period, including the early-2000s downturn, the global financial crisis, and the COVID-19 pandemic. The comparison between asset- and debt-weighted aggregation shows that the concentration of distress among highly indebted firms can differ substantially from the average level of distress in the corporate sector: in 2024 the debt-weighted indicator stood at 31.8%, eight percentage points above the asset-weighted indicator.

Finally, the Shapley decomposition provides an interpretable view of the sources of corporate stress. Across firms and over time, profitability emerges as the main financial dimension underlying predicted distress, accounting for more than four-fifths of the cross-sectional gap, while other categories play smaller and more episode-specific roles.

References

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., Haldeman, R. G., and Narayanan, P. (1977). ZETA analysis: A new model to identify bankruptcy risk of corporations. *Journal of Banking & Finance*, 1(1):29–54.

Altman, E. I., Iwanicz-Drozdowska, M., Laitinen, E. K., and Suvas, A. (2017). Financial distress prediction in an international context: A review and empirical analysis of Altman's Z-score model. *Journal of International Financial Management & Accounting*, 28(2):131–171.

Asquith, P., Gertner, R., and Scharfstein, D. (1994). Anatomy of financial distress: An examination of junk-bond issuers. *The Quarterly Journal of Economics*, 109(3):625–658.

Beaver, W. H. (1966). Financial ratios as predictors of failure. *Journal of Accounting Research*, 4:71–111.

Beck, N., Katz, J. N., and Tucker, R. (1998). Taking time seriously: Time-series–cross-section analysis with a binary dependent variable. *American Journal of Political Science*, 42(4):1260–1288.

Bharath, S. T. and Shumway, T. (2008). Forecasting default with the Merton distance to default model. *The Review of Financial Studies*, 21(3):1339–1369.

Campbell, J. Y., Hilscher, J., and Szilagyi, J. (2008). In search of distress risk. *The Journal of Finance*, 63(6):2899–2939.

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

Chen, K. H. and Shimerda, T. A. (1981). An empirical analysis of useful financial ratios. *Financial Management*, 10(1):51–60.

Duffie, D., Saita, L., and Wang, K. (2007). Multi-period corporate default prediction with stochastic covariates. *Journal of Financial Economics*, 83(3):635–665.

Galán, J. E. (2021). Crews: A camels-based early warning system of systemic risk in the banking sector. Occasional Paper 2132, Banco de España.

Geng, R., Bose, I., and Chen, X. (2015). Prediction of financial distress: An empirical study of listed Chinese companies using data mining. *European Journal of Operational Research*, 241(1):236–247.

Gupta, J., Gregoriou, A., and Ebrahimi, T. (2018). Empirical comparison of hazard models in predicting SMEs failure. *Quantitative Finance*, 18(3):437–466.

Hernandez Tinoco, M. and Wilson, N. (2013). Financial distress and bankruptcy prediction among listed companies using accounting, market and macroeconomic variables. *International Review of Financial Analysis*, 30:394–419.

Lizarzaburu, E. R. (2014). Análisis del modelo z de altman en el mercado peruano. *Universidad & Empresa*, 16(26):137–154.

Manzaneque, M., Priego, A. M., and Merino, E. (2016). Corporate governance effect on financial distress likelihood: Evidence from Spain. *Revista de Contabilidad – Spanish Accounting Review*, 19(1):111–121.

Nam, C. W., Kim, T. S., Park, N. J., and Lee, H. K. (2008). Bankruptcy prediction using a discrete-time duration model incorporating temporal and macroeconomic dependencies. *Journal of Forecasting*, 27(6):493–506.

Ohlson, J. A. (1980). Financial ratios and the probabilistic prediction of bankruptcy. *Journal of Accounting Research*, 18(1):109–131.

Owen, G. (1977). Values of games with a priori unions. In Henn, R. and Moeschlin, O., editors, *Mathematical Economics and Game Theory: Essays in Honor of Oskar Morgenstern*, volume 141 of *Lecture Notes in Economics and Mathematical Systems*, pages 76–88. Springer, Berlin.

Papikova, L. and Papik, M. (2022). Effects of classification, feature selection, and re-sampling methods on bankruptcy prediction of small and medium-sized enterprises. *Intelligent Systems in Accounting, Finance & Management*, 29(4):254–281.

Petersen, M. A. (2009). Estimating standard errors in finance panel data sets: Comparing approaches. *Review of Financial Studies*, 22(1):435–480.

Pham Vo Ninh, B., Do Thanh, T., and Vo Hong, D. (2018). Financial distress and bankruptcy prediction: An appropriate model for listed firms in Vietnam. *Economic Systems*, 42(4):616–624.

Pinches, G. E., Mingo, K. A., and Caruthers, J. K. (1973). The stability of financial patterns in industrial organizations. *The Journal of Finance*, 28(2):389–396.

Pindado, J., Rodrigues, L., and de la Torre, C. (2008). Estimating financial distress likelihood. *Journal of Business Research*, 61(9):995–1003.

Qian, H., Wang, B., Yuan, M., Gao, S., and Song, Y. (2022). Financial distress prediction using a corrected feature selection measure and gradient boosted decision tree. *Expert Systems with Applications*, 190:116202.

Shapley, L. S. (1953). A value for n-person games. In Kuhn, H. W. and Tucker, A. W., editors, *Contributions to the Theory of Games, Volume II*, number 28 in Annals of Mathematics Studies, pages 307–317. Princeton University Press, Princeton, NJ.

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

Taffler, R. J. (1983). The assessment of company solvency and performance using a statistical model. *Accounting and Business Research*, 13(52):295–308.

Tian, S., Yu, Y., and Guo, H. (2015). Variable selection and corporate bankruptcy forecasts. *Journal of Banking & Finance*, 52:89–100.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288.

Vu, N. T., Nguyen, N. H., Tran, T., Le, B. T., and Vo, D. H. (2023). A LASSO-based model for financial distress of the Vietnamese listed firms: Does the COVID-19 pandemic matter? *Cogent Economics & Finance*, 11(1):2210361.

Whitaker, R. B. (1999). The early stages of financial distress. *Journal of Economics and Finance*, 23(2):123–132.

Wruck, K. H. (1990). Financial distress, reorganization, and organizational efficiency. *Journal of Financial Economics*, 27(2):419–444.

Zmijewski, M. E. (1984). Methodological issues related to the estimation of financial distress prediction models. *Journal of Accounting Research*, 22:59–82. Supplement.

A Catalogue of candidate indicators

The table below lists the 39 indicators of the catalogue, grouped by category, with their definition, coverage over the training window and univariate AUROC. Nine indicators belong to the profitability category, six to liquidity, seven to leverage, three to coverage, nine to activity and five to growth and size. Two indicators, the interest coverage ratio and financial debt over EBIT are excluded because they are functions of the EBIT and financial-expense figures that define the distress event in the reserach . Of the remaining 37, eleven fall below the 90% availability floor on the training window, leaving the 26 valid candidates that enter the selection procedure.

Table 14: Catalog of candidate financial indicators

Indicator	Definition	Coverage (%)	AUROC
Profitability (9)			
EBIT / assets (X_3)	EBIT / Total assets	96.7	0.932
ROA (NITA)	Net income / Total assets	96.7	0.875
Retained earnings / assets (X_2)	Retained earnings / Total assets	96.7	0.767
X_2, wide	(Ret. earnings + reserves − treasury shares) / Total assets	96.5	0.741
Net margin	Net income / Sales	90.7	0.855
Operating margin	EBIT / Sales	90.7	0.876
ROE	Net income / Equity (undefined if equity ≤ 0)	96.3	0.877
Earnings momentum (CHIN)	$(NI_t - NI_{t-1})/(\lvert NI_t\rvert + \lvert NI_{t-1}\rvert)$	89.6	0.603
Retained-earnings growth	ΔRet. earnings (wide) / |lag|	89.2	0.691
Liquidity (6)			
Working capital / assets (X_1)	(Current assets − current liab.) / Total assets	96.7	0.628
Current ratio	Current assets / Current liabilities	96.5	0.598
Quick ratio	(Current assets − inventories) / Current liab.	96.5	0.597
Cash ratio	Cash & equivalents / Current liabilities	96.5	0.657
CLCA	Current liabilities / Current assets	96.6	0.596
Working capital / sales	(Current assets − current liab.) / Sales	90.6	0.625
Leverage (7)			
TLTA	Total liabilities / Total assets	96.7	0.522
Equity / debt (X_4, book)	Equity / Total liabilities	96.5	0.523
Long-term debt / assets	Non-current liabilities / Total assets	96.7	0.511
Maturity structure	Current liabilities / Total liabilities	96.5	0.527
Financial debt / EBIT	Financial debt / EBIT (undefined if EBIT ≤ 0)	81.6	0.911
Negative equity (OENEG)	$\mathbf{1}\{\text{Liabilities} > \text{Assets}\}$	96.7	0.500
ΔLeverage	Δ(Liabilities/Assets), year on year	89.5	0.582
Coverage (3)			
Interest coverage	EBIT / |Financial expenses|	96.7	0.962
Average cost of liabilities	|Financial expenses| / Total liabilities	96.5	0.639
Average cost of financial debt	|Financial expenses| / Financial debt	79.2	0.615
Activity (9)			
Asset turnover (X_5)	Sales / Total assets	96.7	0.748
Non-current asset turnover	Sales / Non-current assets	96.5	0.733
Fixed-asset turnover	Sales / Fixed assets (PPE + intangibles + inv. property + subsidiaries)	94.6	0.674
Days sales outstanding	Trade receivables / Sales × 365	90.7	0.583
Inventory turnover	COGS / Inventories	83.9	0.576
Days inventory outstanding	Inventories / COGS × 365	87.2	0.502
Operating expenses / sales	Operating expenses / Sales	90.2	0.718
Days payable outstanding	Trade payables / COGS × 365 (partial coverage)	30.3	0.690
Cash conversion cycle	DSO + DIO − DPO	85.8	0.550
Growth/Size (5)			
Sales growth	ΔSales / |Sales$_{t-1}$|	84.1	0.666
Size	log(Total assets)	96.7	0.626
Age in panel	Years since first appearance	96.7	0.599
Asset growth	ΔAssets / |Assets$_{t-1}$|	89.6	0.627
Fixed-asset intensity	Fixed assets / Total assets	94.6	0.534

Notes: Source: authors' calculations.

B Coefficient estimates for the extended models

Table 15: Estimated coefficients of the expanded logit specifications

	Logit LASSO		Logit parsimonious	
Variable	Coef.	Std. err.	Coef.	Std. err.
Profitability				
EBIT / assets (X_3)	-38.984***	(4.059)	-36.261***	(3.671)
ROA	-1.975	(4.830)	–	
RE / assets (X_2)	2.274	(1.473)	–	
RE / assets (wide)	-1.829	(1.442)	-0.529	(0.578)
Net margin	–		-0.795	(0.620)
Operating margin	-2.262**	(0.974)	-2.515***	(0.932)
ROE	-0.899	(2.263)	–	
Liquidity				
Working capital / assets (X_1)	2.736*	(1.536)	1.059	(1.401)
Quick ratio	-0.290***	(0.103)	-0.477***	(0.131)
Cash ratio	–		0.237	(0.243)
Current liab. / current assets	0.289	(0.329)	0.782**	(0.353)
Working capital / sales	–		0.655	(0.491)
Leverage				
Total liabilities / assets	2.518**	(1.233)	–	
Long-term debt / assets	2.082	(1.951)	–	
Maturity structure	1.343	(0.995)	-0.884	(0.559)
Coverage				
Cost of debt	17.091***	(2.514)	17.156***	(2.349)
Activity				
Asset turnover (X_5)	-1.348**	(0.554)	-1.119***	(0.428)
Non-current asset turnover	-0.331	(0.276)	–	
Fixed-asset turnover	0.190	(0.136)	0.133	(0.100)
Days sales outstanding	0.004**	(0.002)	0.005***	(0.002)
Growth/Size				
Size (log assets)	–		0.037	(0.082)
Firm age	-0.057***	(0.021)	-0.061***	(0.021)
Fixed-asset intensity	0.583	(0.711)	–	
Constant	-2.922***	(1.093)	-1.094	(1.171)
Observations	2,313		2,329	
McFadden R^2	0.539		0.526	
AIC	953.3		991.0	
Predictors	19		16	

Notes: maximum-likelihood logit on the training window (2000–2019); all predictors are lagged one year and winsorised at the 5th and 95th percentiles of the training sample. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Source: authors' calculations.

Table 16: Estimated coefficients of the dynamic specifications

	Panel logit		Discrete-time hazard	
Variable	Coef.	Std. err.	Coef.	Std. err.
Macro-financial controls (lagged one year)				
Credit spread	-0.241***	(0.073)	-0.199**	(0.081)
FX depreciation	0.003	(0.027)	0.006	(0.027)
Real GDP growth	0.049	(0.073)	0.071	(0.074)
Inflation	0.087	(0.095)	0.072	(0.095)
$\ln(1 + \text{age})$	–		-0.812	(0.623)
Real interest rate	0.116*	(0.062)	0.092	(0.067)
Terms of trade	0.016	(0.014)	0.013	(0.014)
Constant	-1.308	(1.797)	-0.727	(1.896)
Observations	2,216		2,216	
McFadden R^2	0.547		0.548	
AIC	889.1		889.6	
Predictors	25		26	

Notes: both specifications add the six macro-financial variables to the LASSO predictor set; the hazard model adds the log of firm age as its baseline term. Estimation on the training window (2000–2019), firm-clustered standard errors in parentheses. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Source: authors' calculations.

C Robustness

C.1 Alternative definitions of the distress event

Table 17 re-estimates the parsimonious logit under three definitions of the event: the baseline coverage criterion , two consecutive years of negative net income, and an interest-coverage threshold of 1.5 instead of 1.

Table 17: Parsimonious logit under alternative definitions of distress

Definition	N	Events	AUROC	Type I	Type II	Cutoff
Coverage, 2 consecutive years (baseline)	581	98	0.969	0.103	0.020	0.139
Negative net income, 2 consecutive years	614	80	0.968	0.101	0.013	0.097
Interest coverage ratio < 1.5	586	117	0.958	0.117	0.094	0.193

Notes: out-of-sample period 2020–2025. Source: authors' calculations.

the results from the alternative definitions suggest that the Parsimonious logit retains its discriminatory power, with a maximum of the Type II error of 10%.

C.2 Alternative classification thresholds

Table 18 reports the parsimonious logit under five threshold rules: the Youden index of Equation (7), the regular $P > 0.5$ rule, and three cost-weighted cutoffs that penalise the Type II error at ratios of 2:1, 5:1 and 10:1.

Table 18: Parsimonious logit under alternative classification thresholds

Rule	Cutoff	Type I	Type II
Youden index	0.139	0.103	0.020
$P > 0.5$	0.500	0.039	0.240
Cost ratio 2:1	0.235	0.074	0.080
Cost ratio 5:1	0.139	0.103	0.020
Cost ratio 10:1	0.135	0.109	0.020

Notes: out-of-sample period 2020–2025.
Source: authors' calculations.

The regular $P > 0.5$ rule shows the cost of ignoring the class imbalance: it cuts the Type I error to 3.9% but misses 24% of distress events. The Youden cutoff, by contrast, keeps the Type II error at 2.0% at the cost of a 10.3% false-alarm rate. More conservative cost ratios do not reduce the Type II error further and only increase false alarms