Early warning detection using Logic Learning Machine: Evidence from private firms

Enrico Ferrari[*], Roberto Garelli[**], Alessandro Limon[***], Alessandro Piazza[****], Lorenzo Simoni[*****], Damiano Verda[******]

Abstract

Purpose: The paper aims to assess the ability of explainable artifiical intelligence (XAI), specifically Logic Learning Machine (LLM), to predict early signals of distress in private companies.

Design/methodology/approach: We examined a sample of Italian private firms that activated early recovery procedures, which are matched to healthy firms. The proprietary algorithm developed by Rulex Innovation Labs is used to discriminate between dis-

[*] Rulex Innovation Labs, Via Felice Romani 9/2, 16122 Genoa (Italy). E-mail: enrico.ferrari@rulex.ai.

[**] Department of Economics and Business Studies, University of Genoa, Via Vivaldi 5, 16126 Genoa (Italy). E-mail: rgarelli@economia.unige.it.

[***] DSTech, Via Salaria 719a, 00138 Roma (Italy). E-mail: a.limon@dstech.it. *At the time the study was developed, Alessandro Limon worked for Rulex Innovation Labs, Via Felice Romani 9/2, 16122 Genoa (Italy).*

[****] Rulex Innovation Labs, Via Felice Romani 9/2, 16122 Genoa (Italy). E-mail: alessandro.piazza@rulex.ai.

[*****] Corresponding author, Department of Economics and Business Studies, University of Genoa, Via Vivaldi 5, 16126 Genoa (Italy). E-mail: lorenzo.simoni@unige.it.

[******] Rulex Innovation Labs, Via Felice Romani 9/2, 16122 Genoa (Italy). E-mail: damiano.verda@rulex.ai.

Conflict of interest disclosure

Roberto Garelli and Lorenzo Simoni have no competing interest to declare.

Enrico Ferrari, Alessandro Piazza, and Damiano Verda are employed at Rulex Innovation Labs, which is the entity that developed and commercialized the algorithm used in the analysis (Logic Learning Machine, LLM).

Alessandro Limon was employed at Rulex Innovation Labs, which is the entity that developed and commercialized the algorithm used in the analysis (Logic Learning Machine, LLM), during the initial stages of this study.

Doi: 10.3280/fr202516015

tressed firms and healthy companies based on a set of publicly available data. Results are then compared with those obtained using other (widely employed) methods.
Findings: The analysis shows that the LLM method is able to classify distressed firms with high accuracy, outperforming logit models and other AI-based methods.
Originality/value: We contribute to the literature on the use of AI in insolvency prediction by exploring the predictive ability of XAI. We also extend the literature on insolvency in private firms, which represent a fundamental part of the economic system and are subject to less scrutiny than public firms.
Practical implications: Our results have practical implications considering the recently enforced EU Insolvency Directive, which imposes the implementation of early warning tools that should be easy to use for all entities across all Member States. By using publicly available data on early distress procedures activated by companies, we build an early warning detection system that can be easily employed by companies of all sizes and types.

Keywords: insolvency, early warning, distress prediction, artificial intelligence, Logic Learning Machine, private firms

JEL: G33, M40

First submission: March 6, 2023. Accepted: October 7, 2024.

1. Introduction

Bankruptcy represents an event that threatens not only a firm and its shareholders but also other stakeholders, such as suppliers, customers, lenders, and other subjects that have relationships with those parties (Kolay et al., 2016; Radovanovic & Haas, 2023). For this reason, the European Union (EU) has revised insolvency legislation, issuing Directive 1023/2019 (Insolvency Directive or Directive afterwards), which introduces early warning procedures to timely detect distress, prevent bankruptcy, and allow companies to restore their financial health when distress occurs. The Directive recommends the use of early warning tools or indicators that should be easy to use for all kinds of companies, allowing them to detect signals of distress at a time when there is still the possibility to put in place recovery procedures, thus surviving the crisis (Bava et al., 2020). This provision, aiming at helping companies survive, is especially important for private firms, which represent a fundamental part of the EU economic system and are subject to less scrutiny than public firms (Mafrolla & D'Amico, 2017; Peek et al., 2010), thus causing greater uncertainty surrounding those companies (Burgstahler et al., 2006; Guedhami & Pittman, 2008).

Several studies have developed insolvency prediction models tailored to private firms, resorting to different methods and employing different triggering events as signals of distress. Among the most frequently used methods, those based on artificial intelligence (AI) play a pivotal role (Varetto, 1999). Research has shown that AI-based methods outperform other methods, such as logistic regression, in different settings (Radovanovic & Haas, 2023).

However, AI solutions used to explore insolvency detection in prior studies, such as support vector machines (SVMs) or neural networks (NNs), suffer from several issues. In particular, it is well known that models such as SVMs and NNs are black boxes; that is, they provide a prediction using complex mathematical formulas that can hardly be understood by humans. Since many efforts are being made both by researchers and by institutions to guarantee the trustworthiness, reliability, and transparency of AI-based models, black boxes should be avoided when dealing with delicate applications. Moreover, black boxes are likely to be little actionable: if the mechanisms leading to a prediction are not clear, it is also difficult to take countermeasures to avoid undesired situations. In other words, black boxes could be very accurate, but often humans do not know what to do with black box predictions. For this reason, a much-debated topic in the past few years has been the so-called explainable AI (XAI), which allows us to overcome some of the issues related to black boxes.

XAI includes AI methods capable of providing predictions that can be explained to human users. This feature allows obtaining results and predictions that can be more easily understood and interpreted by users. Logic Learning Machine (LLM) is an XAI solution that has been employed in several research studies in different fields, demonstrating high performance and reliability (Ferrari et al., 2023). Interestingly, no scientific works have employed LLM to predict early warning signals and compare the results obtained with those derived from other AI-based methods.

Against this background, this study employs the LLM method to forecast early signals of distress in a sample of Italian private companies based in the three most important regions for Gross Domestic Product (GDP), namely Lombardy, Lazio, and Veneto. In particular, the analysis was conducted using a proprietary algorithm developed by Rulex Innovation Labs.[1] The algorithm identifies the combination of variables that best predicts the outcome of analyzing large data sets. It also formulates some rules that can be employed to predict potential distress situations.

[1] The algorithm was not developed for this specific study but had already been developed and employed in several settings before.

The results are then compared with those obtained through logistic regression and other AI-based methods widely used in the literature, namely SVMs, NNs, and decision trees (DTs), to understand whether XAI can offer more valuable and reliable estimates regarding early distress prediction.

The focus on Italy allows us to control for the influence of the local economic context, which has been indicated as a factor capable of influencing distress prediction. Some studies show that models tailored to a specific context yield better results (Dainelli et al., 2013). To take this aspect into account, we followed prior research and restricted the analyses to a single country.

Italy has implemented an early warning procedure defined by law, anticipating the EU Insolvency Directive. In its initial form, the early warning procedure was based on the identification of distress using an analysis of indicators derived from financial statements (Bava et al., 2020). The use of financial statement data to forecast insolvency was conceived as an easy-to-implement tool, as all Italian companies in the form of limited liability entities must publish financial statements regardless of their size or sector. The public availability of financial statement data in Italy also makes it an ideal setting to gain insights into insolvency among private entities.

In contrast to studies using bankruptcy as a triggering event, which does not enable predicting early distress signals (Balcaen & Ooghe, 2006), we considered the procedures activated by companies to recover when early signs of distress are detected (Altman et al., 2009; Kalay et al., 2007). Unlike prior studies employing this kind of signal, however, we focus on private firms. In Italy, companies can resort to some procedures that allow them to partially repay creditors and restructure debt composition through an agreement. In this study, we use those procedures, as information about their activation is publicly available. The aim of composition with creditors and debt restructuring procedures is to allow an entity to avoid bankruptcy and survive a (temporary) situation of potential distress.

As early warning tools should be easy to understand and employ by all types of organizations in the view of EU regulators, we resorted to the most common financial variables employed by the academic literature and/or identified by professional bodies, combined with non-financial indicators retrievable from publicly available sources for private entities, to predict insolvency.

The findings reveal that the LLM method accurately detects early signs of distress, outperforming logit models, SVMs, NNs, and DTs. Hence, this study contributes to the literature in several ways. First, it extends previous studies on AI applications to insolvency prediction by examining the performance of XAI in the form of LLM. While the LLM method (Rulex algorithm) has been used in other fields, this is the first research study to employ this method in insolvency prediction.

Second, we contribute to the development of a distress prediction model that is easy to implement and tailored to a specific context. Specifically, by restricting our analysis to Italy, we can consider the specific characteristics of companies operating in that context. Third, we expand the literature on insolvency prediction among private companies by considering the role of early signals of distress.

The study also has practical implications. Following the recent regulatory reforms, Italian companies, like entities located in other EU countries, must predict signals of early distress to timely activate a recovery procedure if needed. Thus, our results can be of interest to companies that have to monitor potential signs of distress. In line with the recommendations of the EU and the Insolvency Code, an early warning model based on financial statement data and publicly available non-financial information could be a viable and easy-to-implement solution, even for small entities.

2. Background

Recent regulatory changes in the EU emphasized the need for all companies, including small and medium enterprises, to monitor their financial health and timely detect signals of potential distress. Those regulatory changes are connected to the approval of EU Directive 1023/2019, which aims to harmonize the characteristics of insolvency regulations in the European area. The Directive states that all Member States must adopt some early warning tools that can timely detect signals of distress. The goal of this provision is to allow distressed companies to start a recovery process, thus avoiding bankruptcy. The capability to recover from a distress situation largely depends on capturing early signals of a potential crisis. According to the European regulator, early warning tools should be able to detect distress situations when they are still reversible so that entities can timely activate some recovery processes.

Despite the request to define early warning tools, the Directive does not specify which characteristics those tools should have. According to the Directive, such tools could include alert mechanisms that signal that a debtor has not made certain types of payments or that some indicators suggest potential distress, but they could also include advisory services provided by public or private organizations. Among the most employed tools to predict distress, insolvency prediction models play a pivotal role. Those tools usually employ a set of variables, which can be either financial data or other types of data (e.g., data about ESG, ownership, and industry), to discriminate between distressed and healthy companies based on a sample that includes both types of firms.

Accordingly, some Member States might adopt early warning mechanisms based on insolvency prediction models. For example, Italy enforced a new Code of Insolvency (approved with Legislative Decree 14/2019), anticipating the Insolvency Directive, according to which firms are required to identify signals of distress considering specific financial indicators. The new Code of Insolvency issued in Italy established a new procedure that aims to offer distressed companies assistance by an independent, qualified expert in the definition of a feasible recovery plan. The procedure must be activated when early distress signals emerge. The recognition of those signals is based on early warning tools specifically designed to address this issue. The development of early warning tools featured several steps. Initially, the national legislator appointed the National Association of Chartered Accountants, which is called the Consiglio Nazionale dei Dottori Commercialisti e degli Esperti Contabili (CNDCEC), to develop an early warning tool. The CNDCEC elaborated a set of financial indicators that could be calculated based on a company's financial statements. The model was developed by an independent agency with experience in the development of insolvency prediction models.[2] According to the model, every company had to monitor those indicators, and, if certain thresholds were exceeded, the board of directors and the supervisory board (if present) had to signal the potential state of distress to an authority, which then should appoint an independent expert to assist the entity[3].

To develop an early warning instrument that could be easily adopted by all kinds of entities, including small entities that are not equipped with an adequate information system that allows them to collect and elaborate complex data or to elaborate prospective financial information, the CNDCEC re-

[2] The agency is CERVED, an entity that manages, elaborates, and distributes financial statement data and other data about Italian entities.

[3] The procedure was articulated in different steps, requiring the evaluation of different indicators. First, a company had to consider the shareholders' equity. In case the equity was negative, the entity must activate the recovery procedure, without the necessity of further assessments. In case the equity is positive, the early warning tool required the company to evaluate the ratio of the sum of current cash holdings and future cash inflows to future cash outflows related to debt reimbursement. If the ratio is equal to or higher than 1, it indicates that the company can face debt reimbursement. If the ratio was lower than 1, the company must activate the recovery procedure. A particular characteristic of this ratio is that, according to the CNDCEC, companies should calculate pro forma cash inflows and outflows based on some projections. In the absence of a reliable estimate of prospective data, the entity should consider five indicators calculated based on the financial statements (see footnote 4).

sorted to a set of financial ratios. The ratios identified by the CNDCEC are based on the dominant literature in the fields of financial statement analysis and insolvency prediction and are indicators of profitability, liquidity, financial structure, and interest coverage[4].

The most recent amendments to the Code of Insolvency (Legislative Decree 83/2022) led to the abandonment of this set of indicators in favor of a single indicator, which is the prospective debt service coverage ratio (DSCR). This indicator, calculated as the ratio of the future cash inflows to the future cash outflows related to debt reimbursement, indicates that a company can face debt reimbursement in the period considered if the ratio is equal to or higher than 1. According to the Code of Insolvency, a crisis manifests itself as the inadequacy of prospective cash flows to cover debt reimbursement in the following 12 months. At the same time, a company must be capable of elaborating a plan that illustrates whether and how it will be able to reasonably repay its debt in the future. That plan is the basis for determining whether the company can access the recovery plan. These changes to the original early warning tool were made immediately before the enforcement of the new Code of Insolvency in the summer of 2022.[5] While the new approach has some advantages, being even simpler than the early warning tool based on financial indicators, the requirement to use prospective data could result in high costs for some entities, particularly small companies. However, as the objective of the Code of Insolvency is to foster the adoption of internal control systems that can timely detect crises, insolvency prediction models could still be useful for companies to monitor their financial health and the viability of the business. Having insolvency prediction tools that are easy to implement has become a priority for all the entities based in Italy after the enforcement of the new Code. Potentially, all companies in the EU area will be interested in easy-to-use insolvency prediction tools given the transposition of the EU Insolvency Directive by the Member States.

[4] The indicators were the following:
- the ratio of interests to revenues;
- the ratio of equity to debt;
- the ratio of current assets to current liabilities;
- cash return on assets (ROA), measured as the ratio of cash flows to total assets;
- the ratio of debt toward the fiscal authority to total assets.

[5] The Code should have been enforced in 2020, but due to the COVID-19 pandemic, it has become effective since July 2022.

3. Prior literature and research purpose

3.1. Insolvency prediction using AI

Several scholars have employed AI-based methods to predict firm distress. These methods emerged as a valid alternative to models based on linear regression analysis, as some research studies highlighted that in many cases the assumptions on which linear discriminant models are based are not supported by empirical evidence. AI allows us to overcome some of the issues related to linear model assumptions. Moreover, empirical studies have demonstrated that, in several cases, those models have better performance than logistic regression models (Radovanovic & Haas, 2023). The use of AI in insolvency prediction became popular in the 1990s (D'Annunzio & Falavigna, 2004; Varetto, 1999).

Due to their characteristics, models based on DTs and NNs have been widely employed to build insolvency prediction models. DTs allow classifying companies in a sample into different groups based on similar characteristics. NNs are programmed by an analyst to learn how to distinguish insolvent companies from healthy companies, considering specific characteristics that are used as inputs mutually interconnected through a series of nodes, each with a specific weight (Brockett et al., 1994; Chung et al., 2008; Tsai & Wu, 2008).

However, most of the studies employing AI in bankruptcy prediction analyze large corporations. Crisis detection in private firms remains a relatively underexplored area, as only a few studies have focused on default prediction in private companies (Matenda et al., 2022), often resorting to logit models (Charalambakis & Garrett, 2019) or adaptations of the *z*-score (Altman et al., 2017; Jacoby et al., 2019; Range et al., 2018). An example of an AI-based application to insolvency prediction is the study by Papana and Spyridou (2020), who analyzed a set of private firms in Greece using bankruptcy as the triggering event. The authors compared DTs, NNs, discriminant analysis, and logit regression. Other studies focused on contexts in which (small and medium) private entities provide a substantial contribution to the economy, but they also included large firms in the analysis (e.g., Moscatelli et al., 2019). Recent studies have also incorporated natural language information, such as notes to the financial statements or narrative information disclosed in the annual report, processed through deep learning techniques (Borchert, 2023), to improve the accuracy of the prediction. A systematic survey of techniques used for the prediction of business distress is provided by Scherger et al. (2019).

Some research studies on AI and insolvency focused on specific contexts, such as the study by Varetto (1998), who analyzed Italian entities in the period between 1982 and 1995, comparing linear discriminant analysis and genetic algorithms. The author shows that genetic algorithms are very effective in detecting distress, but, in some cases, linear discriminant analysis was found to be better. A more recent study conducted by Moscatelli et al. (2019) on a sample of Italian companies found that AI-based methods are more effective in predicting default than linear discriminant models. The authors conducted a series of additional analyses, showing that the superior predictive capability of AI is reduced when (a) some non-financial variables on the historical credit rating of a firm are added and (b) small samples are analyzed.

A potential issue that emerges in the above-mentioned studies is connected to the triggering event used to identify distressed companies. Bankruptcy is the event most frequently used as the triggering event. However, bankruptcy often corresponds to the final phase of a crisis, which can be seen as a process (Balcaen & Ooghe, 2006). Hence, those studies could fail to assess early signs of distress.

3.2. The use of XAI and the role of LLM in AI applications

Current research trends in AI emphasize the importance of trustworthiness and reliability for the development of (semi)automatic decision-making systems. The trustworthiness of AI-based solutions has been assessed in different studies and is also supported by public institutions such as the European Commission. For example, Ala-Pietilä et al. (2020) identify seven requirements that need to be satisfied to ensure the trustworthiness of AI solutions: human agency, robustness, data governance, transparency, fairness, societal and environmental well-being, and accountability. Moreover, the study highlights how the possibility for a human to understand the decisions taken by an AI-based method is crucial to ensuring the fulfillment of these requirements. Explainability has become one of the cross-cutting key topics to broaden the applicability and acceptability of AI-based solutions. Explainability could be approached in different ways depending on the requirements and the application fields, either focusing on local explainability, which provides an explanation for a single instance, or global explainability, which explains all the possible cases. For example, a local explanation would explain why distress is likely to occur for a specific firm, while a global explanation would explain why firms in general are likely to go into distress. While local explanations could be very useful for analyzing a single prediction made by AI, global explanations are more effective in improving the general knowledge about the studied phenomenon. DTs are probably the most well-known method for AI explanation. Nonetheless, classifiers based on DTs are known to be weak; that is, their accuracy is poor. Moreover, the explanation provided by dichotomous rules is often limited because the rules are usually built by dividing the input space into very small subsets with a divide-and-conquer approach. LLM is a classification algorithm based on a Switching Neural Network (SNN) model (Muselli, 2006), which can overcome these drawbacks thanks to some good properties that will be explained in the next sections. The LLM model is currently used in some fintech-related fields, such as anti-fraud, credit scoring, and non-performing loan (NPL) performance predictions.

Besides these applications, the LLM model has been proven to provide good-quality predictions and explanations in cross-cutting application fields such as diagnosis support (Gerussi et al., 2022; Parodi et al., 2015; Verda et al., 2019) and complex system optimization (Ferrari et al., 2023).

In particular, comparative studies have highlighted that LLM obtains better results than DTs and that the results are comparable to those of other well-known methods, such as SVMs and NNs (Muselli, 2012). The rules obtained by LLM are usually richer and more general than DT rules, allowing a better understanding of the mechanisms leading to a prediction. Moreover, thanks to their structure, LLM rules are more actionable; that is, it is possible (for a human or an automated system) to use the information contained in the rules to make decisions that can improve the quality of the system (Ferrari et al., 2023).

3.3. LLM and early distress prediction

Considering the above-mentioned background, the study aims to employ XAI, in particular the LLM method, to predict early distress situations among private firms, comparing this method with other methods commonly employed in distress prediction, such as logit and AI-based methods, to understand whether XAI can outperform traditional AI methods in this field.

The study also aims to test the predictive ability of publicly available data in the detection of early signals of distress among private firms. While early signs of distress have been employed in studies examining large, listed firms (Altman et al., 2009; Kalay et al., 2007), we use them as early warning signs in private firms, thus extending the literature on distress in these kinds of entities. Insolvency prediction in private firms represents an underexplored field, which could be enriched by adopting novel methods or exploring novel settings.

Identifying factors that enable the prediction of distress could help develop an early warning tool that companies can adopt. Hence, we build an insolvency prediction model that is easy to use by any company, in line with what is required by the EU Insolvency Directive and the new Code of Insolvency enforced in Italy.

4. Methodology

4.1. Logic Learning Machine

We used the LLM method, as implemented in the Rulex software (www.rulex.ai), to analyze the data. LLM is an optimized implementation of a model named a Switching Neural Network (SNN), whose functioning has been explained in previous scientific articles (Muselli, 2006). The objective of the analysis is to investigate the capability of different sets of variables to forecast early distress. An SNN or LLM is a classification method that allows the automatic generation of a model directly from a set of labeled data. Unlike other similar methods, such as NNs and SVMs, LLM creates a model based on IF-THEN rules, which can therefore be understood by an experienced user. Thanks to this feature, LLM falls within the paradigm of XAI, which includes AI methods capable of providing predictions that can be explained to human users. To generate the rules, LLM converts the original variables to a Boolean domain, thus transforming each example of the original table into a string of 0s and 1s. Starting from the transformed table, it is possible to adopt appropriate techniques for the synthesis of Boolean functions, such as the shadow clustering algorithm (Muselli & Ferrari, 2011). The Boolean functions reconstructed in this phase are then transformed into as many sets of IF-THEN rules that are needed for the final classification model. Thanks to appropriate algorithms for synthesizing Boolean functions, the rule sets are constructed in such a way as to maximize generality (i.e., the number of cases that each rule describes) and simplicity (i.e., the number of conditions contained in each rule). To measure the generality of a rule, two quantities are usually defined:

- covering, defined by TP/(TP + FN), which measures the percentage of cases correctly described by the rule, where TP = true positive and FN = false negative;
- error, defined by FP/(FP + TN), which measures the percentage of cases incorrectly described by the rule, where FP = false positive and TN = true negative.

A small error is typically allowed to maximize the generality of the rules, especially for data sets with a strong presence of noise. Figure 1 shows a schematic representation of the Switching Network model, as explained in Muselli (2006). In the first phase (latticization), each variable is transformed into a string of binary data using the inverse only-one code binarization, and all strings are eventually concatenated into one unique large string per subject. In the second phase, the shadow clustering algorithm (Muselli & Ferrari, 2011) is used to generate a set of binary vectors (the implicants), each of which identifies a cluster in the input space associated with a specific output class. Finally, all the implicants are transformed into simple conditions and combined into a set of intelligible rules.

Figure 1 – Schematic representation of the Logic Learning Machine (LLM) algorithm

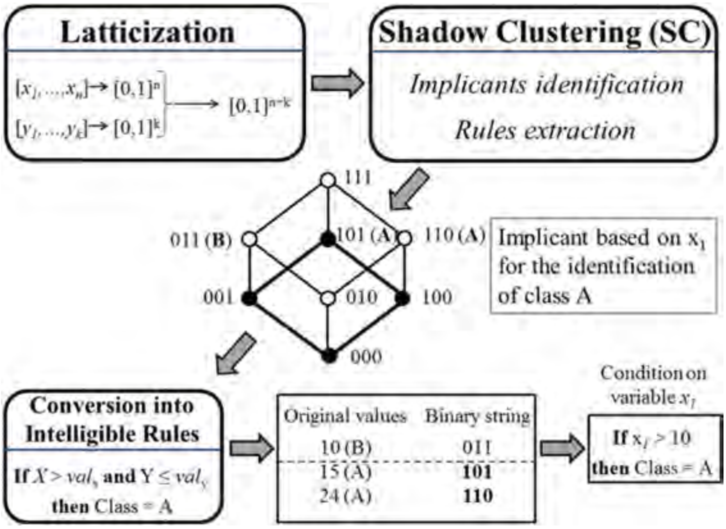

Source: Muselli (2006)

A characteristic of the rules generated by an SNN or LLM is their overlap. In contrast to divide-and-conquer methods, such as DTs, LLM constructs rules using the entire data set. Therefore, it is possible for an example to be covered by more than one rule. The cases described by several rules of different classes (in this case, a healthy company or a state of difficulty) are typically intermediate situations; thanks to the overlapping rules, the researchers can build different classifiers that are more or less unbalanced toward the two classes.

4.2. Sample

Prior studies have emphasized how the socioeconomic context might influence distress prediction. For instance, models developed using large, listed companies based in the United States could be non-optimal in predicting distress in different samples composed of firms operating in different geographical areas or of different sizes. For this reason, some scholars have developed distress prediction models tailored to specific types of companies (e.g., Altman & Sabato, 2007; Dainelli et al., 2013) or specific geographical areas (Giacosa et al., 2015). We focused on Italy for three main reasons. First, private firms play a pivotal role in the Italian economy and GDP. Second, financial statement data and data about the activation of early distress recovery procedures are publicly available for private firms. Third, Italy was among the first countries to reform insolvency regulation to foster the application of early warning indicators by defining a specific alert mechanism by law.

We selected our sample by considering companies operating in the three most important Italian regions for their contribution to GDP, according to the Italian Bureau of Statistics (Istituto Nazionale di Statistica [ISTAT]). The three regions are Lombardy, Lazio, and Veneto. The focus on different regions allows us to capture differences in regional economic growth (Brasili et al., 2012), as well as in other economic factors, such as employment and labor conditions (Culotta et al., 2022) or the resilience of an area to economic shocks due to the presence of certain industries (Lagravinese, 2015). We started with non-financial firms found in the Aida Bureau van Dijk (BvD) database based in the selected regions with early recovery procedures (*concordato preventivo*) or debt restructuring agreements that occurred between 2013 and 2019. We selected this time frame to be sure that our data were not affected by the global financial crisis, the COVID-19 pandemic, or the Ukraine war. As we aimed to employ one-year-lagged measures to forecast distress, we focused on events that occurred between 2014 and 2019. If a company had more than one event within this time frame, we considered only the first event that occurred. Once companies were identified, we removed firms without financial data for the year immediately before the event.

We then built a control sample by means of propensity score matching. Non-financial firms based in the same three regions that were active at the date of data extraction and that had available financial statements for all the years between 2013 and 2019 were selected. We considered only industries with both companies with procedures (distressed companies) and companies with no signs of distress. For each industry identified by the ATECO 2-digit code, we matched each distressed firm with five healthy companies using propensity score matching based on one-year-lagged total assets and a caliper of 0.01. One-year-lagged total assets referred to the year when the procedure began for distressed companies and a reference year randomly selected between 2014 and 2019 for the other companies. Due to the impossibility of finding five matches for each observation using the selected caliper and the removal of companies with negative age according to the Aida BvD database, the final sample features 432 distressed companies (showing either an early recovery procedure or a debt restructuring agreement) and 2,129 matched healthy companies.

4.3. Training and out-of-sample prediction

The standard approach in model extraction and evaluation includes the definition of training and out-of-sample sets. In this approach, the training set is used to extract the model, and the out-of-sample set is then used to apply the model and evaluate the performance robustness on new data. Usually, 70% of the available data, randomly selected, constitutes the training set, while the remaining 30% constitutes the out-of-sample set. Another popular choice for the split is 80% training data and 20% out-of-sample data.

To guarantee that the observed performance is not due to a particularly lucky or unlucky split, leading to an overestimation or underestimation of the performance, respectively, we adopt the well-known k-fold approach (Gareth et al., 2013). In other words, we follow the following procedure:

1. The data set is split into k groups of equal size (in the experiments, k = 10).
2. For k times, the model is extracted using all the groups, except the kth group, as a training set and then applied to the kth group, which constitutes the test set for that run.
3. The performance on the test set is measured as the average performance on the test set across the k runs. Similarly, the performance on the training set is measured as the average performance on the training set across the k runs.

This allows us to mitigate the dependence of the performance on the split, considering that each row belongs to the test set in one run and to the training set in k - 1 runs, thus equalizing the contribution of each prediction to the performance on the training and the test set.

4.4. Variables

4.4.1. Dependent variable

The dependent variable used to predict a crisis is a dichotomous variable that takes a value of 1 if early seizure signs are found and 0 otherwise. While most previous studies used bankruptcy as the triggering event, this choice can prevent detecting early distress signals. Bankruptcy is usually the final step of a long process that causes the deterioration of a firm's financial health. Predicting bankruptcy might not be sufficient to avoid distress, depending on the length of the period of crisis and the time lag used to predict the event (Balcaen & Ooghe, 2006). While catching the crisis as it arises is one of the main challenges of insolvency forecasting studies, some researchers have overcome this issue using the activation of procedures aimed at restoring businesses as an early warning signal (Altman et al., 2009; Kalay et al., 2007).

In line with those studies, a variable was used that indicates the activation of procedures aimed at allowing the continuation of business activity after overcoming a moment of difficulty. In particular, two procedures provided by Italian law were considered in this study: the composition with creditors and the debt restructuring agreement. The first procedure consists of finding an agreement with some of the creditors of the potentially insolvent company to satisfy some of them through business continuity or the liquidation of the company's assets. A prerequisite for accessing this procedure is the existence of a state of crisis. However, this procedure is often used to allow the company to continue operating, signaling a state of temporary and remediable crisis. Therefore, the composition with creditors can be considered a usable tool before the crisis becomes irreversible, leading to the bankruptcy of the company.

The second procedure is the debt restructuring agreement. This procedure is aimed at reducing debt exposure by rebalancing the financial structure. It provides for the renegotiation of the amounts to be repaid and the financing conditions with creditors, requiring the approval of many of them. Creditors who do not participate must be fully satisfied. Companies that activate this procedure experience difficulties in repaying the debt, thus showing signs of crisis, albeit slight in some cases. Therefore, such a procedure can be considered an early sign of crisis. From this perspective, the above-mentioned procedures can be considered similar to the Chapter 11 procedure in force in the United States (Altman et al., 2009; Kalay et al., 2007). Following prior studies, we built a dummy variable that takes a value of 1 if a company activated one of these procedures in a certain year and 0 otherwise (crisis).

4.4.2. Variables used to predict distress

We used a set of financial and non-financial indicators that can capture a state of economic or financial difficulty based on prior studies. An advantage of using the LLM method is that it allows the formulation of some rules using combinations of the variables that better predict the outcome. Hence, the algorithm automatically selects the variables that, in combination, predict distress better.

In line with the literature on insolvency prediction in private firms (Matenda et al., 2022) and the study developed by the CNDCEC on the occasion of the first version of the new Italian Insolvency Code, we included profitability ratios, variables that consider the financial structure, and coverage ratios. Specifically, we considered the return on assets (roa) (Charalambakis & Garrett, 2019; Jacoby et al., 2019), the return on sales (ros) (Slefendorfas, 2016), the return on equity (roe) (Slefendorfas, 2016), and cash return on assets (cash_roa) (CNDCEC, 2019) as profitability measures. We used the ratio of equity to total assets (financial_ind) as a measure of leverage (Altman et al., 2017; El Khoury & Al Beaïno, 2014). Interest coverage, measured as the ratio of earnings before interests, expenses, depreciation, and amortization to interest expenses, was used as a coverage measure (coverage) (Comacho-Minano et al., 2015). Following CNDCEC (2019), we included the ratio of interest expenses to revenues (int_rev). We also included raw measures for total debt (debt) and interest expenses (int_exp).

As negative equity is considered an important predictor of distress (Orlando & Rodano, 2020), we included the book value of equity (equity) and a dummy variable that takes a value of 1 if a company has negative equity and 0 otherwise (neg_equity). We also considered total revenues (revenues) (Jones & Wang, 2019); operating margins (Jones & Wang, 2019), namely earnings before interests and taxes (ebit) and earnings before interests, taxes, depreciation, and amortization (ebitda); cash flow from operations (cash_flow); and unscaled net income (earn) (Camacho-Minano et al., 2015; Jones & Wang, 2019). In line with prior studies, we included retained earnings (ret_earn), the ratio of retained earnings to total assets (ret_earn_assets) (Barboza et al., 2017; Radovanovic & Haas, 2023), capital turnover (turnover) (Range et al., 2018), working capital (wc) (Jones & Wang, 2019), and working capital scaled by total assets (wc_assets) (Barboza et al., 2017; Radovanovic & Haas, 2023). All independent financial variables are lagged by one year. If a company had revenues equal to 0, we replaced the ratio of interest expenses to revenues with the highest determinable value in the sample and return on sales with the lowest determinable value in the sample. If the interest expenses were equal to 0, coverage was replaced with the highest determinable value in the sample.

We also considered non-financial variables. We included the age of the company (age) (Camacho-Minano et al., 2015), measured as the number of years from its establishment; the one-year lagged number of employees (empl); and the type of company, indicating whether a company is an artisan firm (artisan) or an innovative small or medium-sized enterprise (SME) (inn_sme). Details on all the variables employed are presented in Table 1.

Table 1 – Variable definition

Variable	Description	Time lag
Dependent variable		
crisis	A dummy variable that takes a value of 1 if a company entered into a composition with creditors procedure or a debt restructuring agreement, 0 otherwise	t (no time lag)
Independent financial variables		
cash_flow	Cash flow for the year	t-1
cash_roa	The ratio between cash flow and total assets	t-1
coverage	The ratio between EBITDA and interest expenses	t-1
debt	Total debt	t-1
earn	Net income for the year	t-1
ebit	Earnings before interests and taxes	t-1
ebitda	Earnings before interests, taxes, depreciation, and amortization	t-1
equity	Book value of equity	t-1
financial_ind	The ratio between equity and total assets	t-1
int_exp	Interest expenses accrued	t-1
int_rev	The ratio between interest expenses and revenues	t-1
neg_equity	A dummy variable that takes a value of 1 if equity is negative, 0 otherwise	t-1
ret_earn	Retained earnings	t-1
ret_earn_assets	Ration between retained earnings and total assets	t-1
revenues	Revenues for the year	t-1
roa	Ratio between operating profit and total assets	t-1
roe	Ratio between net profit and equity	t-1
ros	Ratio between operating profit and sales	t-1
turnover	Capital turnover, determined as the ratio between revenues and total assets	t-1
wc	Working capital	t-1
wc_assets	The ratio between working capital and total assets	t-1
Independent non-financial variables		
age	Number of years since the establishment of the company	t (no time lag)
artisan	A dummy variable that takes a value of 1 if a company is labeled as "artisan" in the Bureau van Djik databank	t (no time lag)
empl	Number of employees	t-1
inn_sme	A dummy variable that takes a value of 1 if a company is labeled as an "innovative SME" in the Bureau van Djik databank	t (no time lag)

Descriptive statistics (Table 2) show that 15% of the companies show a negative value of equity. Firms in the sample are aged 24 years on average.

Table 2 – Descriptive statistics for the dependent and financial independent variables (full sample)

Variable	Mean	Q1	Median	Q3	St. Dev.
crisis	0.17	0	0	1	0.37
cash_flow	247.44	-22.64	85.14	446.91	9384.7
cash_roa	-0.03	-0.01	0.02	0.06	0.61
coverage	3820.83	0.09	3.79	19.59	151680.5
debt	15505.56	1370.87	3965.49	11535.08	72555.39
earn	-900.85	-103.79	10.68	162.22	18153.9
ebit	-202.37	-58.96	67.19	339.164	10979.49
ebitda	925.58	-0.80	155.91	622.86	9763.9
equity	7369.99	113.8	986.11	4432	55556.17
financial_ind	29.25	7.26	23.59	49.028	30.15
int_exp	313.04	5.84	39.19	152.93	3089
int_rev	5.79	0.01	0.013	0.05	194.27
neg_equity	0.15	0	0	1	0.35
ret_earn	-467.86	0	0	0	9537.61
ret_earn_assets	-0.94	0	0	0	34.37
revenues	15495.78	436.31	2521.41	7961.73	106209.2
roa	-0.03	-0.01	0.02	0.05	0.37
roe	0.25	-0.001	0.047	0.22	8.45
ros	-7.37	-0.034	0.033	0.1	255.44
turnover	0.79	0.072	0.65	1.19	0.91
wc	11895.9	1066.87	3339.57	9004.36	45208.34
wc_assets	0.62	0.37	0.69	0.91	0.32

Notes: cash_flow, debt, earn, ebit, ebitda, equity, int_exp, ret_earn, revenues, and wc are expressed in €/thousands.

The correlation matrix (heatmap in Figure 2) shows that some of the variables employed seem to be highly correlated. For example, int_rev and ros have a correlation coefficient of 0.998, and earn and ebit have a correlation coefficient of 0.914, highlighting the fact that for these two couples, the two variables contain similar information. Notice that for standard linear methods, this could be an issue, since the presence of collinear inputs could worsen the performance of the classifier. Nonetheless, using LLM allows us to overcome this issue since one of the collinear variables is selected to be included in the model, without affecting its quality.

Figure 2 – Correlation matrix between numerical attributes considered

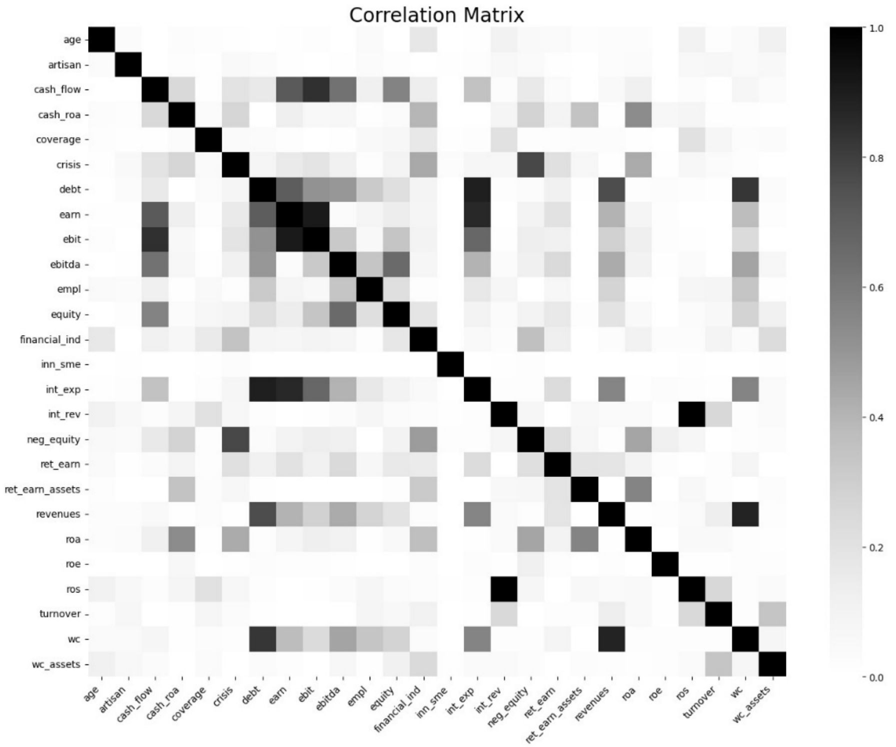

5. Results

5.2. Insolvency prediction using LLM

Figure 3 shows the five variables that the model considers the most important ones for its prediction and to what extent each variable is relevant: the longer the histogram bar associated with it, the more relevant the variable is.

Figure 3 – The ranking of features involved in the model

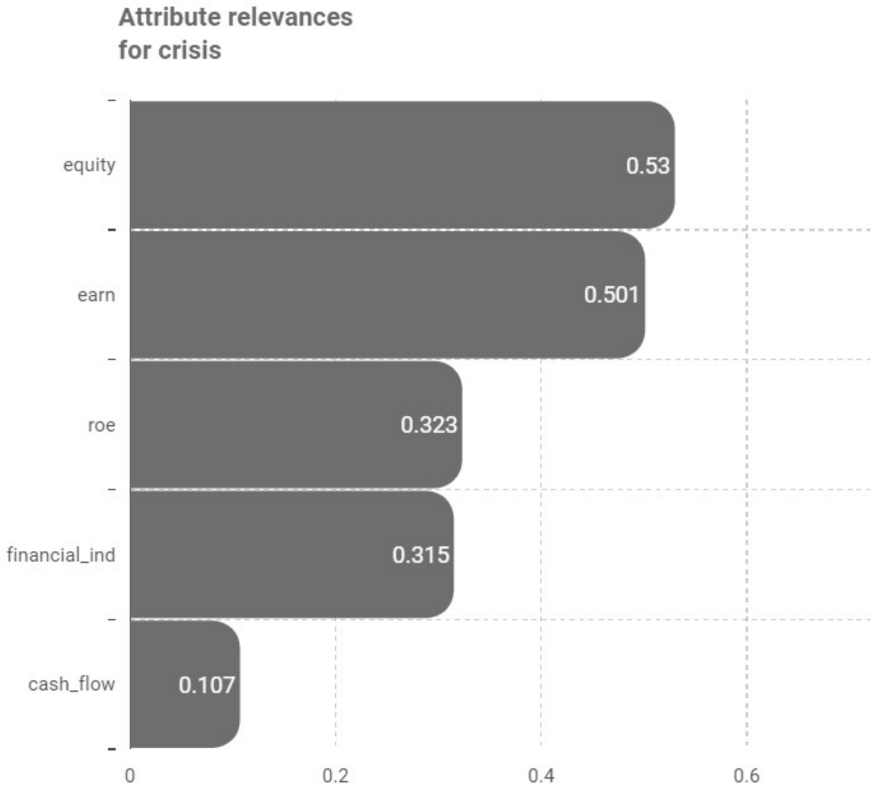

The most important variables are equity, net income (earn), return on equity (roe), leverage (financial_ind), and cash flow of the period (cash_flow).

The results confirm the predictive ability of some of the factors that are most frequently used in the literature. For instance, leverage has been extensively included in studies on private firms, demonstrating to be a significant factor in predicting distress, as more leveraged companies are more likely to go bankrupt (e.g., Charalambakis & Garrett, 2019; Jones & Wang, 2019). Our findings are in line with those studies and demonstrate that the financial independence of the company is also a predictor of early distress, proxied by the use of recovery procedures.

Several studies have revealed that profitability indicators can predict corporate distress. In particular, our evidence confirms previous findings related to the predictive ability of return on equity (Jones & Wang, 2019) in detecting early distress.

The relevance attributed by the algorithm depends not only on the direct correlation between the predictive variable and the dependent variable but also on the contribution that the predictive variable, in combination with the other predictive variables, makes to the prediction of the dependent variable.

These variables are combined by the model in the form of rules, such as the following one:

Rule 1*: IF equity < -13.40 THEN "crisis"*

The rule highlights the fact that companies with negative equity with a value below 13.4 thousand euros will tend to be subject to a crisis in the following year. Another rule extracted by the LLM is the following:

Rule 2: *IF financial_ind < 43.49 AND ret_earn_assets <= -0.024 AND turnover <= 1.57 THEN "crisis"*

The second rule highlights that companies with a ratio of equity to total assets, retained earnings on total assets, and capital turnover below certain thresholds are more at risk of crisis than others.

Figure 4 summarizes the effectiveness of the model in predicting the dependent variable. The upper left cell represents the number of cases in which "no crisis" is correctly predicted, and the upper right cell represents the cases in which "no crisis" is not correctly predicted. In the bottom row, we find in the cell on the left the number of cases in which "crisis" is predicted as "no crisis", while the cell on the right shows the number of cases in which "crisis" is predicted correctly. As we can see, if we concatenate all the test sets of the 10-fold cross-validation procedure (i.e., if we consider the prediction on each sample without having trained the model on that sample), the proposed model correctly identifies 88.2% of the crisis cases and 92.7% of the non-crisis cases. Only in 7.3% of the non-crisis cases (156 out of 2,129) does the model register a false positive; that is, it predicts a crisis anyway.

Figure 4 – The confusion matrix of the predictions made by the LLM model on the concatenation of the test sets, in 10-fold cross-validation

Confusion Matrix

targetAttribute	Forecast	0	1
Actual	***2561***	*2129*	*432*
0	*2024*	1973 (92.673%)	156 (7.327%)
1	*537*	51 (11.806%)	381 (88.194%)

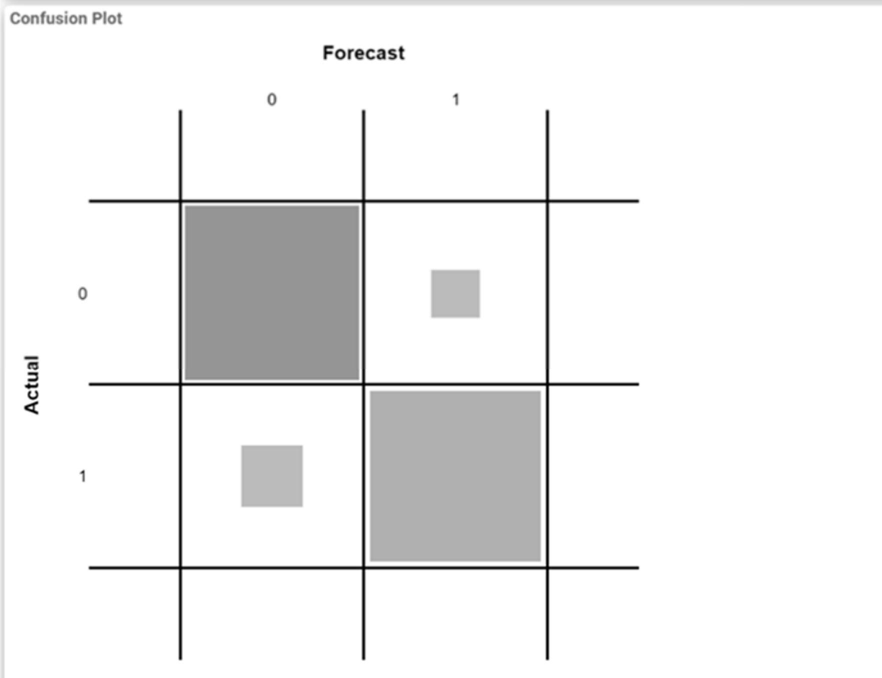

As can be seen in Table 3, this 7.3% of cases are concentrated in a small number of ATECO codes. In this case, to achieve a more granular analysis of the results, we considered the ATECO six-digit code. Even though we have a high number of distinct values when using the six-digit code (more specifically, 539 sectors), we notice that only 66 sectors include at least one false positive and, even more significantly, that only 4 sectors include at least 10 false positives (Table 3).

Table 3 – Sectors of activity comprising at least 10 false positives

Sector	False Positives (n.)
Construction of residential and non-residential buildings	33
Development of real estate projects without construction	21
Purchase and sale of immovable property carried out on own property	17
Real estate leasing	10

In addition to the binary forecast (“crisis”/“no crisis”), the model also provides a score, namely a numerical value between 0 and 1, which represents the probability of the crisis occurring. Depending on the threshold chosen to distinguish the “crisis” outputs from the “no crisis” outputs, different classifiers are defined. The confusion matrix in Figure 4 is obtained by applying one of the possible thresholds (0.5), which represents the threshold value associated with the classifier that maximizes the value of the Youden index (Youden, 1950).

Considering this continuous score, we can calculate another performance metric, concerning all possible binarization thresholds. This metric is called the AUC (area under the curve), and it is equivalent to the area under the receiver operating characteristic (ROC) curve, as shown in Figure 5. As we can see, the AUC is, in this case, equal to 0.965 (on the concatenation of all the test sets). This means that by randomly choosing a company in crisis and a company not in crisis, there is a 96.5% probability that the model will assign the company in crisis a higher score (probability of crisis) than the other company.

Figure 5 – ROC Curve obtained by changing the threshold in the score to determine the crisis output, on the concatenation of the test sets, in 10-fold cross-validation

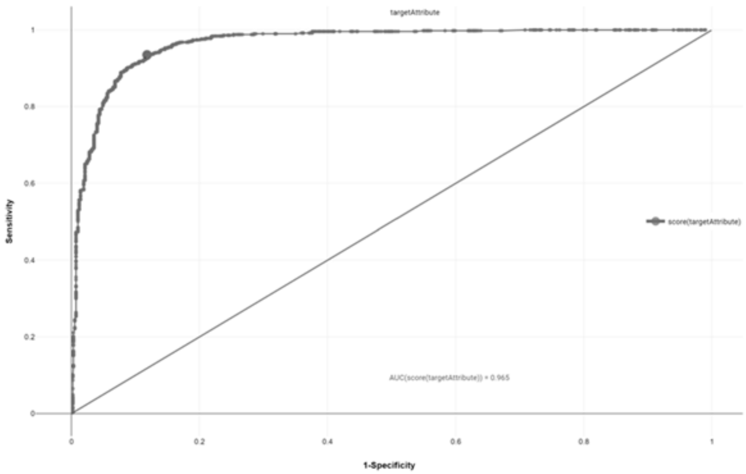

Overall, our distress prediction model can forecast an early warning procedure with high accuracy, thus providing a tool that can help understand the path toward a potential insolvency situation. The model highlights that a combination of some financial variables allows distinguishing companies that will enter a composition with creditors or a debt restructuring agreement.

5.3. Comparison between the LLM method and other methods

To evaluate the performance of the LLM method more robustly, we apply other statistical and machine learning techniques to the same problem, evaluating their performance with the same 10-fold cross-validation procedure described in Section 4.3.

We employed logistic regression, an SVM, an NN, and a DT. Logistic regression is a very simple method that uses a linear combination of inputs to compute a score function that is used (according to a threshold) to discriminate between positive and negative cases. An SVM employs linear optimization techniques to find the classifier in an augmented space that reduces the misclassification while keeping the highest separation between positive and negative cases. In the experiments, a radial basis function was used with regularization parameter $C = 0$. An NN is a very popular method based on several perceptrons (i.e., simple linear classifiers) organized in feedforward layers. The first layer is the input layer, the last layer provides the output, and the intermediate layers are the so-called hidden layers. In these experiments, one hidden layer with five perceptrons was used. DTs are popular classifiers based on a sequence of tests that determines the path followed by a sample within the tree until it reaches a leaf, which is associated with a decision (positive, negative). Thanks to their structure, DTs are considered clear boxes. In fact, the tree can be converted into a set of IF-THEN rules. In the experiments, the tree was grown up to pure leaf nodes and then pruned with an entropy-based method.

Our results (Table 4) show that the performance of the LLM method on the test set is superior to the performance obtained using logistic regression and the SVM, NN, and DT methods. The LLM method has the highest AUC, as well as the highest covering (the ratio of true positives to the sum of true positives and false negatives). Overall, the LLM method can provide reliable insights into early warning signals for private entities.

Table 4 – Comparison between LLM, Logistic Regression, SVM, Neural Network, and Decision Trees

Model	True Negative	False Negative	False Positive	True Positive	AUC
LLM	1973	51	156	381	0.965
Logit	2026	150	103	282	0.931
SVM	2082	110	47	322	0.776
NN	2086	111	43	321	0.958
DT	2110	113	19	319	0.865

6. Conclusions

Recent regulatory changes have emphasized the role of early distress prediction and early warning systems. While early warning tools have not been defined in detail by the EU, Italy has adopted a new insolvency regulation in which some indicators calculated based on financial statement data were initially defined as early warning tools. Despite those indicators being subsequently dismissed, the importance of timely prediction of distress remains an important issue for Italian companies subject to the new Insolvency Code.

Against this background, we have examined a sample of Italian private firms, resorting to LLM to develop a model capable of predicting early warning distress. We show that XAI, in the form of the LLM algorithm employed, can forecast insolvency in private firms with a high level of accuracy, outperforming other methods commonly employed in insolvency prediction.

Our study contributes to the literature in different ways. First, we extend the literature on insolvency prediction in private firms by documenting the role of the LLM method. Prior studies using AI in insolvency prediction relied on different methods. While LLM has been applied in other fields, to the best of our knowledge this is the first study using LLM in insolvency prediction.

Second, we examine early distress in private firms by employing composition with creditors and debt restructuring agreements as signals of early distress instead of bankruptcy, which has been criticized for representing the final step of a lengthy deterioration process (Balcan & Ooghe, 2006). While prior studies have employed such events (Altman et al., 2009; Kalay et al., 2007), we extend the use of those signals to capture early warning distress in private firms.

The study has also practical implications, as we developed an early warning tool that can be easily employed by all types of companies, adhering to the purpose of the new EU Insolvency Directive and the Italian Insolvency

Code enforced in 2022. According to the EU Insolvency Directive, early warning tools should be easy to use and employable by all kinds of entities. As Italian firms are required to prepare financial statements, they can easily draw from those documents to elaborate the ratios we employed to identify distressed firms. The recent EU reform on insolvency requires companies to evaluate their ability to continue to operate in the foreseeable future to timely activate a recovery procedure. Hence, the model we have developed could be of interest to companies that need an internal early default prediction system that is easy to implement.

Finally, our study has implications for policymakers who are involved in an ongoing process aimed at defining early warning tools. We show the contribution that XAI and the algorithm that we used could make to the development of an easy-to-use early distress prediction model.

References

Ala-Pietilä, P., Bonnet, Y., Bergmann, U., Bielikova, M., Bonefeld-Dahl, C., Bauer, W., Bouarfa, L., Chatila, R., Coeckelbergh, M., Dignum, V., & Gagné, J.F. (2020). *The assessment list for trustworthy artificial intelligence (ALTAI)*. European Commission.

Altman, E. I., & Sabato, G. (2005). Effects of the New Basel Capital Accord on Bank Capital Requirements for SMEs. *Journal of Financial Services Research*, 28, 15-42. Doi: 10.1007/s10693-005-4355-5.

Altman, E. I., Kant, T., & Rattanaruengyot, T. (2009). Post-Chapter 11 Bankruptcy Performance: Avoiding Chapter 22. *Journal of Applied Corporate Finance,* 21(3), 53-64. Doi: 10.1111/j.1745-6622.2009.00239.x.

Altman, E. I., Iwanicz-Drozdowska, M., Laitinen, E. K., & 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. Doi: 10.1111/jifm.12053.

Balcaen, S., & Ooghe, H. (2006). 35 years of studies on business failure: an overview of the classic statistical methodologies and their related problems. *The British Accounting Review*, *38*(1), 63-93. Doi: 10.1016/j.bar.2005.09.001.

Barboza, F., Kimura, H., & Altman, E. (2017). Machine learning models and bankruptcy prediction. *Expert Systems with Applications*, 83, 405-417. Doi: 10.1145/3374549.3374550.

Bava, F., Di Trana, M. G., & Cane, M. (2020). Can a quantitative approach be mitigated? Proposals for the application of the "early warnings" required by the new Italian Insolvency Code. *Financial reporting,* 2, 33-61.

Borchert, P., Coussement, K., De Caigny, A., & De Weerdt, J. (2023). Extending business failure prediction models with textual website content using deep learning. *European Journal of Operational Research*, 306(1), 348-357. Doi: 10.1016/j.ejor.2022.06.060.

Brasili, C., Bruno, F., & Saguatti, A. (2012). Economic growth and dualism in Italian regions: A spatiotemporal model. *Rivista Italiana degli Economisti*, 3, 397-416.

Brockett, P. L., Cooper, W. W., Golden, L. L., & Pitaktong, U. (1994). A neural network method for obtaining an early warning of insurer insolvency. *Journal of Risk and Insurance*, 61, 402-402. Doi: 10.2307/253568.

Burgstahler, D. C., Hail, L., & Leuz, C. (2006). The importance of reporting incentives: Earnings management in European private and public firms. *The Accounting Review*, 81(5), 983-1016. Doi: 10.2308/accr.2006.81.5.983.

Camacho-Miñano, M. D. M., Segovia-Vargas, M. J., & Pascual-Ezama, D. (2015). Which characteristics predict the survival of insolvent firms? An SME reorganization prediction model. *Journal of Small Business Management*, 53(2), 340-354. Doi: 10.1111/jsbm.12076.

Charalambakis, E. C., & Garrett, I. (2019). On corporate financial distress prediction: What can we learn from private firms in a developing economy? Evidence from Greece. *Review of Quantitative Finance and Accounting*, 52(2), 467-491. Doi: 10.1007/s11156-018-0716-7.

Chung, K. C., Tan, S. S., & Holdsworth, D. K. (2008). Insolvency prediction model using multivariate discriminant analysis and artificial neural network for the finance industry in New Zealand. *International Journal of Business and Management*, 39(1), 19-28. Doi: 10.5539/ijbm.v3n1p19.

Consiglio Nazionale dei Dottori Commercialisti e degli Esperti Contabili (CNDCEC) (2019). *Crisi d'Impresa. Gli Indici dell'Allerta*. Consiglio Nazionale dei Dottori Commercialisti e degli Esperti Contabili (CNDCEC).

Culotta, F., Alaimo, L. S., Bravo, J. M., di Bella, E., & Gandullia, L. (2022). *Total-employed longevity gap, pension fairness and public finance: Evidence from one of the oldest regions in EU. Socio-Economic Planning Sciences*, 82, Part A, 101221. Doi: 10.1016/j.seps.2021.101221.

Dainelli, F., Giunta, F., & Cipollini, F. (2013). Determinants of SME credit worthiness under Basel rules: the value of credit history information. *PSL Quarterly Review*, 66(264), 21-47.

D'Annunzio, N., & Falavigna, G. (2004). *Modelli di analisi e previsione del rischio di insolvenza: una prospettiva delle metodologie applicate*. Ceris-Cnr.

El Khoury, R., & Al Beaïno, R. (2014). Classifying manufacturing firms in Lebanon: An application of Altman's model. *Procedia-Social and Behavioral Sciences*, 109(1), 11-18.

European Union Directive 2019/1023 of the European Parliament and of the Council of 20 June 2019 on preventive restructuring frameworks, on discharge of debt and disqualifications, and on measures to increase the efficiency of procedures concerning restructuring, insolvency and discharge of debt (2019). -- https://eur-lex.europa.eu/eli/dir/2019/1023/oj.

Ferrari, E., Verda, D., Pinna, N., & Muselli, M. (2023). Optimizing Water Distribution through Explainable AI and Rule-Based Control. *Computers*, 12(6), 123. Doi: 10.3390/computers12060123.

Gerussi, A., Verda, D., Cappadona, C., Cristoferi, L. Bernasconi, D.P., Bottaro, S., Carbone, M. Muselli, M., Invernizzi, P., & Asselta, R. (2022). LLM-PBC: Logic Learning Machine-Based Explainable Rules Accurately Stratify the Genetic Risk of Primary Biliary Cholangitis. *Journal of Personalized Medicine*, 12, 1587. Doi: 10.3390/jpm12101587.

Giacosa, E., Mazzoleni, A., Teodori, C., & Veneziani, M. (2015). Insolvency prediction in companies: an empirical study in Italy. *Corporate Ownership and Control*, 12(4), 232-350. Doi: 10.22495/cocv12i4c1p6.

Guedhami, O., & Pittman, J. (2008). The importance of IRS monitoring to debt pricing in private firms. *Journal of Financial Economics*, 90(1), 38-58. Doi: 10.1016/j.jfineco.2007.12.002.

Italian Government Legislative Decree 12 January 2019, n. 14 (2019). -- https://www.gazzettaufficiale.it/dettaglio/codici/codiceCrisi.

Italian Government Legislative Decree 17 June 2022, n. 83 (2022). -- https://www.gazzettaufficiale.it/eli/id/2022/07/01/22G00090/sg.

Jacoby, G., Li, J., & Liu, M. (2019). Financial distress, political affiliation and earnings management: the case of politically affiliated private firms. *The European Journal of Finance*, 25(6), 508-523. Doi: 10.1080/1351847X.2016.1233126.

Jones, S., & Wang, T. (2019). Predicting private company failure: A multi-class analysis. *Journal of International Financial Markets, Institutions and Money*, 61, 161-188. Doi: 10.1016/j.intfin.2019.03.004.

Kalay, A., Singhal, R., & Tashjian, E. (2007). Is Chapter 11 costly?. *Journal of Financial Economics*, 84, 772-796. Doi: 10.1016/j.jfineco.2006.04.001.

Kolay, M., Lemmon, M., & Tashjian, E. (2016). Spreading the misery? Sources of bankruptcy spillover in the supply chain. *Journal of Financial and Quantitative Analysis*, 51(6), 1955-1990. Doi: 10.1017/S0022109016000855.

Lagravinese, R. (2015). Economic crisis and rising gaps North-South: evidence from the Italian regions. *Cambridge Journal of Regions, Economy and Society*, 8(2), 331-342. Doi: 10.1093/cjres/rsv006.

Mafrolla, E., & D'Amico, E. (2017). Borrowing capacity and earnings management: An analysis of private loans in private firms. *Journal of Accounting and Public Policy*, 36(4), 284-301. Doi: 10.1016/j.jaccpubpol.2017.05.001

Matenda, F. R., Sibanda, M., Chikodza, E., & Gumbo, V. (2022). Bankruptcy prediction for private firms in developing economies: a scoping review and guidance for future research. *Management Review Quarterly*, 72, 927-966. Doi: 10.1007/s11301-021-00216-x.

Moscatelli, M., Narizzano, S., Parlapiano, F., & Viggiano, G. (2019). *Corporate default forecasting with machine learning*. Banca d'Italia working papers, n. 1256.

Muselli, M. (2006). *Switching Neural Networks: A New Connectionist Model for Classification*. Neural Nets – WIRN NAIS 2005 2005. Lecture Notes in Computer Science, vol 3931, Springer.

Muselli, M. (2012). Extracting knowledge from biomedical data through Logic Learning Machines and Rulex. *EMBnet.journal*, 18, 56-58. Doi: 10.14806/ej.18.B.549.

Muselli M., & Ferrari E. (2011). Coupling logical analysis of data and shadow clustering for partially defined positive Boolean function reconstruction. *IEEE Transaction on Knowledge and Data Engineering*, 23, 37-50. Doi: 10.1109/tkde.2009.206.

Orlando, T., & Rodano, G. (2020). *Firm undercapitalization in Italy: business crisis and survival before and after COVID-19*. Banca d'Italia occasional papers, n. 590.

Papana, A., & Spyridou, A. (2020). Bankruptcy prediction: the case of the Greek market. *Forecasting*, 2(4), 505-525. Doi: 10.3390/forecast2040027.

Parodi, S., Filiberti, R., Marroni, P., Libener, R., Ivaldi, G.P., Mussap, M., Ferrari, E., Manneschi, C., Montani, E., & Muselli, M. (2015). Differential diagnosis of pleural mesothelioma using Logic Learning Machine. *BMC bioinformatics*, 16, 1-10. Doi: 10.1186/1471-2105-16-S9-S3.

Peek, E., Cuijpers, R., & Buijink, W. (2010). Creditors' and shareholders' reporting demands in public versus private firms: Evidence from Europe. *Contemporary Accounting Research*, 27(1), 49-91. Doi: 10.1111/j.1911-3846.2010.01001.x.

Radovanovic, J., & Haas, C. (2023). The evaluation of bankruptcy prediction models based on socio-economic costs. *Expert Systems with Applications*, 227, 120275. Doi: 10.1016/j.eswa.2023.120275.

Range, M. M., Njeru, A., & Waititu, G. A. (2018). Using Altman's Z score (Sales/Total Assets) Ratio Model in Assessing Likelihood of Bankruptcy for Sugar Companies in Kenya. *International Journal of Academic Research in Business and Social Sciences*, 8(6), 683-703.

Scherger, V., Terceño, A., & Vigier, H. (2019). A systematic overview of the prediction of business failure. *International Journal of Technology, Policy and Management*, 19(2), 196-211. Doi: 10.1504/IJTPM.2019.100601.

Šlefendorfas, G. (2016). Bankruptcy prediction model for private limited companies of Lithuania. *Ekonomika*, 95(1), 134-152.

Tsai, C. F., & Wu, J. W. (2008). Using neural network ensembles for bankruptcy prediction and credit scoring. *Expert Systems with Applications*, 34(4), 2639-2649. Doi: 10.1016/j.eswa.2007.05.019.

Vaertto, F. (1998). Genetic algorithms applications in the analysis of insolvency risk. *Journal of Banking & Finance*, 22, 1421-1439. Doi: 10.1016/S0378-4266(98)00059-4.

Varetto, F. (1999). Metodi di previsione delle insolvenze: un'analisi comparata. In G. Szego & F. Varetto (Eds.), *Il Rischio Creditizio: Misura e Controllo* (pp. 178-301). Torino: Utet.

Verda, D., Parodi, S., Ferrari, E., & Muselli, M. (2019). Analyzing gene expression data for pediatric and adult cancer diagnosis using logic learning machine and standard supervised methods. *BMC bioinformatics*, 20(9), 1-13. Doi: 10.1186/s12859-019-2953-8.