Decision Support Systems 107 (2018) 64–77

Contents lists available at ScienceDirect

Decision Support Systems

journal homepage: www.elsevier.com/locate/dss

Failure pattern-based ensembles applied to bankruptcy forecasting

Philippe du Jardin

Edhec Business School, 393 Promenade des Anglais, BP 3116, Nice Cedex 3 06202, France

ARTICLE INFO

Article history:
Received 27 August 2017
Received in revised form 16 January 2018
Accepted 17 January 2018

Keywords:
Decision support systems
Ensemble-based models
Self-organizing map
Bankruptcy forecasting

ABSTRACT

Bankruptcy prediction models that rely on ensemble techniques have been studied in depth over the last 20 years. Within most studies that have been performed on this topic, it appears that any ensemble-based model often achieves better results than those estimated with a single model designed using the base classifier of the ensemble, but it is not uncommon that the results of the former model do not outperform those of a single model when estimated with any other classifier. Indeed, an ensemble of decision trees is almost always more accurate than a single tree but not necessarily more than a neural network or a support vector machine. We know that the accuracy of an ensemble used to forecast firm bankruptcy is closely related to its ability to capture the variety of bankruptcy situations. But the fact that it may not be more efficient than a single model suggests that current techniques used to handle such a variety are not completely satisfactory. This is why we have looked for a method that makes it possible to better embody this diversity than current ones do. The technique proposed in this article relies on the quantification, using Kohonen maps, of temporal patterns that characterized the financial health of a set of companies, and on the use of an ensemble of incremental size maps to make forecasts. The results show that such models lead to better predictions than those that can be achieved with traditional methods.

1. Introduction

Models that have been studied in the financial literature and that are used to forecast bankruptcy are primarily default models: a firm goes bankrupt when it lacks sufficient resources to meet its financial obligations, hence when it becomes insolvent. Most empirical studies that have focused on bankruptcy prediction have therefore attempted to find measures that characterize a risk of default. The first models developed in the 1960s, following the study by Altman [4], have sought to assess this risk by estimating the distance between the financial situation of a given firm and a standard bankruptcy situation. Virtually all data-mining techniques that have been developed for classification purposes have been used to design failure models that share almost all the same characteristics: models are dichotomous, have good forecasting abilities and are easy to estimate. However, what can be considered the main factor of their success is also their main weakness. They essentially rely on a single rule and are estimated using financial data that solely characterize a unique period of firm life. This type of modeling reflects a rather rudimentary view of bankruptcy; it is considered the result of a a-historical process [39] that does not depend on time and that is reducible to a limited number of measures. But reality is a bit different. One knows that firms that apparently share the same financial profile, from the point of view of a model, may in reality have a very different probability of failure. Over time, some of them may have gained a certain resilience that gives them the ability to withstand failure. Some others may have received from their environment a sort of carrying capacity that has changed their fate at the very moment where their situation worsened, or have managed to recover even though nothing suggested they were able to do so [11]. All these factors, which can solely be analyzed over time, cannot be properly embodied by traditional models.

The historical dimension of failure and the multiplicity of the situations that lead to bankruptcy have given rise to a large body of literature. We can find, on the one side, studies that focused on the temporal dimension of financial failure. They analyzed the way variables that measure firm activity over several years [22] may influence model accuracy, assuming that taking time into account with multi-period data would be sufficient to embody the dynamics of the phenomenon. We can also find, on the other side, studies that were interested in modeling the different financial situations that lead to bankruptcy. They especially analyzed how to embody at-risk situations using ensemble-based models, this time assuming that the multiplication of forecasting rules would make it possible to model the diversity of failure symptoms [32,46,61]. In both cases, models on the whole lead to better results than those estimated with single models, but it seems that multi-rule models sound more promising

E-mail address: philippe.dujardin@edhec.edu.

https://doi.org/10.1016/j.dss.2018.01.003

than multi-period ones. Indeed, the sole use of historical data hardly changes model physiognomy; they are still singular and still stumble on the fact that they embody a unique measure of a distance to bankruptcy which is too simple to be truly effective. In fact, ensemble models make it possible to represent the different facets of the decision boundary that separate failed from non-failed firms [9,49] and, to a certain extent, are able to embody the different patterns of decline that may lead to failure and that traditional models cannot assess.

For all that, when one looks closely at the results achieved with all these models and compares those estimated with ensemble techniques to those calculated with the best single models, in many cases, discrepancy between them does exist but it is relatively low (between 2 and 3%) and is often is not statistically significant [21,23,27,43,49,63]. Of course, a model designed with an ensemble of decision trees is often more accurate than one estimated with a single tree, but not necessarily much more powerful than a single model estimated with a support vector machine. This shows that ensembles represent a source of performance that is related to the way they are able to capture the variety of bankruptcy situations. But the fact that ensembles are not systematically more accurate than single models suggests that they do not handle variety in a completely optimal way. These are the reasons why we have studied a method that makes it possible to better embody the different bankruptcy situations, using what we call "failure patterns", and to use these "patterns" to make forecasts. It relies on the quantification, with Kohonen maps, of temporal patterns that characterize both failed and non-failed firms, and on the use of an ensemble of incremental size maps to make forecasts.

2. Literature review

The very first bankruptcy prediction models that were developed, following that of Altman [4], represent failure as if it were a critical situation that might be captured using a measure of the discrepancy that separates the financial situation of a given firm from that of a standard critical situation. These models also make the assumption that failure is a phenomenon that can be estimated using a unique measure of firm financial health. Finally, they assume that failure corresponds to a particular event that can be explained using a single classification rule, and to a certain extent that bankruptcy is the result of a unique process of decline. The limits of these models are the direct consequences of their assumptions. We know that bankruptcy has multiple causes and symptoms, and that a model with variables that are solely measured over a single period would probably not be able to embody such diversity. We also know that failure does not depend on the sole situation of a firm at a given period of its life, but is the result a protracted process that cannot be captured properly by models that do not take into account a temporal dimension. Finally, we know that different paths to firm failure exist [11,39,48], but their complexity cannot be properly assessed with simple single-rule models.

In order to overcome these limitations, some research works have sought to estimate and use the historical dimension of bankruptcy through multi-period data, but still using traditional single-rule models. They have shown that models designed with financial variables measured over several years lead to better predictions than those achieved with models designed with single period variables [22,65]. But their structure can solely solve a part of the issue: the uniqueness of the rule leads to models that still embody a single standard bankruptcy situation. This is why some other research works have attempted to model the variety of bankruptcy situations using multiple classification rules. The techniques that have been used for this purpose aim at designing a meta-model where each component has a particular expertise on a certain region of the decision space. If these components are sufficiently diverse [38], they make it possible to estimate models that are on average more accurate than single models.

The characteristics of ensemble-based techniques rely on the way classification rules are developed and combined. With certain techniques, rules can be estimated with the same modeling method. This is the case of bagging, boosting, random subspace… where all models are estimated using either a decision tree [23], logistic regression [43], a feed-forward neural network [32], a survival model [16], k-nearest neighbors [49] or a support vector machine [61]. Rules can also be computed with different methods: for example, a model estimated with a logistic regression can be used in conjunction with other types of models that are assessed either with a support vector machine [25], or with both a support vector machine and a neural network [70], or with a neural network combined with a decision tree and discriminant analysis [46]. Sometimes, combinations are more complex; a first set of models is computed with different methods and their results are then used as inputs of a final model estimated with a neural network [7]. Finally, rules can be estimated after a prior segmentation of the decision space. The technique consists of grouping observations into a few classes and then calculating models where each of them fits the characteristics of a given class, using the same methods as those presented above [15,16,63].

On the whole, ensemble models present a better ability to make accurate forecasts than single models do. However, the absolute gain brought by these techniques is relatively low compared to that of traditional models, even if it remains real. We have measured the average gain calculated over 31 studies published between 2000 and 2017, and presented in Table 1[1]. If one calculates the difference, for each study, between the correct classification rate of the best ensemble-based model and that of the best single model from among those that have been estimated while taking into account the sample size used for their validation[2], the average gain does not even reach 2.4%. And, above all, if one estimates whether the differences between these models are significant, one can notice that among the 31 differences calculated in Table 1, only 10 are statistically significant.

The literature shows that the ability of a model to capture the whole variety of bankruptcy situations is a key factor of its performance. But it also shows that usual ensemble-based models are not able to easily embody this variety because, each rule, solely represents a boundary between two groups, although we might likely consider more subtle modeling. Literature has long shown the existence of different profiles or failure "patterns" which represent prototype situations that firms may experience over their life, and where some of them may lead to bankruptcy. But it has mainly focused on very general patterns that are shared by the largest number of firms. Those that were estimated by D'Aveni [11], Laitinen [39] or by Lukason et al. [48] illustrate this finding. However one may think that these patterns are far too general and could be refined. A widespread cause of bankruptcy, such as a lack of liquidity, can be embodied in many different financial situations that are very unequally distributed within a population of firms. A firm may not be liquid because of a lack of cash, a lack of permanent capital, a problem of balance between payables and receivables…This suggests that non-liquid firms can be represented through a wide variety of profiles and that a large number of illiquidity patterns should

[1] Table 1 lists the main studies that have been published since 2000 and that have studied the accuracy of ensemble-based models when it comes to forecasting bankruptcy. For each study, this table presents the results achieved with the best ensemble-based models (maximum 3) and those achieved with the best single model from among all single models that have been estimated.

[2] When the size of the test sample was not indicated, we estimated the gain using the size of the learning sample, despite the positive bias introduced in the estimation.