ERASMUS UNIVERSITY ROTTERDAM

Erasmus School of Economics

Master Thesis Business & Economics: Financial Economics

Predicting and characterizing the recovery of zombie firms over time

Job Lagewaard

476702jl

Supervisor: Prof. Dr. Tim Eisert

Second assessor: TBA

18-11-2022

Abstract

To investigate what measures can be taken by zombie firms to recover their firms, this research uses several XGBoost models to determine the main characteristics of recovering firms for each year between 2009 and 2022. These results have been used to find their practical implications. It has been found that zombie firms should downsize their workforce, sell fixed assets, and minimize rental expenses while refocusing on the intangible assets already owned by the firm. Moreover, goodwill is found to be increasing in importance, suggesting consolidations to be a tool for zombie firms to recover. Further evidence indicates that discontinuing operations and focusing on increasing sales does not contribute much to the recovery of zombie firms.

Contents

Introduction

Japan, Europe, and the US have been struggling with so-called ‘zombie firms’. A firm becomes a zombie when can afford all necessary operating costs and interests but is not capable of servicing its debts. These firms are artificially kept alive by their lenders because it is less expensive for zombie lenders to keep zombie firms alive indefinitely than to write off the loans. This creates an economically inefficient situation that spills over into the rest of the zombie’s sector and the overall economy (Acharya et al., 2022).

The rationale behind the increasing number of zombie firms during certain periods is often cited to macro-economic shocks. An example of this can be found in the low interest rates in recent years; low interest rates make it relatively easy for zombie firms to keep lending and to stay afloat. However, interest rates started to rise in 2022, these firms will increasingly rely on their zombie lenders, which lowers their chance of survival. Zombie firms can also be caused by policy makers, which was the case in Japan where the government and banks ignored economic problems and wrongly implemented countermeasures to help the economy (Giannetti & Somonov, 2013) (Gaballero, Hoshi & Kashyap, 2008).

Several papers have been written about the conundrum of what the definition of a zombie firm is, and how it affects its direct and indirect economic surroundings. Acharya et al. (2022) found for research focusing on these economic effects, defining zombie lenders can best be done using interest rate subsidies, as it helps to determine the effect of zombie firms on credit misallocation without being dependent on direct firm specifics. Interest rate subsidies are defined as interest rates lower than the market’s cost of capital for healthy firms.

Their paper is the basis for another paper by De Martiis et al. (2021). Using machine learning techniques, De Martiis et al. (2021) tried to predict future zombies in Europe and the US. The model they used is around 90% accurate when it comes to classifying and predicting zombie firms. With it they found a firm’s total assets and annual earnings to be most important when searching for zombie firms, while debt indicating features only influence US-based zombies. Their research shortly taps into the most important features of recovering zombies, finding common stock to be the main driver of a firm's recovery.

It is important to get a better understanding of zombie firms as zombie firms are associated with declining employment growths and sinking investments for non-zombie firms in industries where zombies are present. (Adalet McGowan et al., 2018). Adalat McGowan et al. (2018) argue that the presence of zombie firms in a certain sector limits the possibility to expand for healthy firms and heightens barriers of entering the sector. This is expected to come from the fact that competing against a zombie firm, which is being subsidized, forces newcomers to be extremely efficient before they can expect any profits (Caballero et al., 2008).

Nakamura and Fukuda (2013) examined the persistently elevated levels of Japanese zombie firms between 1990 and 2010 and found, although most zombie firms managed to recover during the 2000s, their effect on Japan's problematic deflation persisted several years after. Using multinomial log regressions, they found selling fixed assets and downsizing workforces to be the best way to develop a firm from its zombie status back to a healthy firm. They further reason the lack of innovation in recovering firms in combination with restructuring workforces causes ongoing deflation. (Nakamura & Fukuda, 2013)

Another paper by Fukuda and Nakamura (2011) also investigated recovered zombies and found Japanese zombie firms hardly ever go bankrupt. Their results indicate the same as the paper above; downsizing and selling fixed assets are the main drivers of a firm's recovery. However, they also found the absence of positive incentives for managers to have a negative effect on the chances of survival (Fukuda & Nakamura, 2011).

These papers show zombie firms have a negative external effect on healthy firms in the same sector. One would expect recovering firms to counter this effect, but the aforementioned papers find these firms disturb inflation and may even cause long-term deflation, which economists commonly dislike.

To get a better understanding of this effect, recovered firms must be investigated. Recovering firms are interesting because existing literature focuses on zombie firms, a quantitative study on recovering firms is rare and mostly based on logistic regressions. From an academic perspective it is important to use new methods to quantify recovering firms and their features as it elaborates our knowledge about them whilst verifying, bringing nuance, or adding to papers based on traditional methods.

Most research focusing on recovering firms has been investigating the most important features of recovered firms. To my knowledge, no paper can be found focusing on characterizing and predicting recoveries over time. This is a rather interesting phenomenon as the prediction of recoveries helps to get a more complete view of what makes a recovered firm a recovered firm.

Gaining insights on recovering firms will further help policy makers make decisions focused on limiting the number of zombie firms and their effect on the overall economy. Further, existing literature elaborates on the negative spillovers caused by recovering firms, the creation of a model that to a greater extent maps the characteristics of recovering firms and the drivers of said recoveries will help policy makers and managers to get a better grip on the negative spillovers, making it possible to take more effective countermeasures.

There is a gap in existing literature when it comes to research focused on recovering firms: firstly because of a lacking variety in used models and secondly because no significant research has been focusing on the characteristics of firm recoveries over time.

The first gap will be filled with the use of machine learning models that allow for the many different inputs. The use of such models allows for an investigation with many more firm-specific features than traditional methods can cope with, which helps to find patterns that may seem implausible or unlikely for the human mind.

For the second gap, insights will be gained on what can be done by managers to resurrect a zombie firm based in Europe or North America. Public firms active between 2004 and 2022 are used to characterize zombie firms, but more importantly, their recovery. By balancing the datasets towards recovering firms, in combination with a more versatile machine learning model than can be found in existing literature (XGBoost), it is possible to find the most important factors of a firm's recovery over time.

The findings stemming from XGBoost indicate the retained earnings of a zombie firm are the most important factor for its recovery. This very intuitive result is trivial at first glance, but upon further research, it must be concluded that the increase in retained earnings for recovering firms cannot come from investments in the company, as zombie firms lack the capability of doing these investments. Increasing one's retained earnings via profit margins was also found to be implausible because zombie firms are not able to increase their prices. The only way recovering firms can increase their retained earnings is by cutting costs.

The results further indicate cost cutting should not come from lowering dividends, discontinuing operations, or trying to decrease fixed costs per unit by focusing on sales. On the contrary, cost cutting should be done by lowering rental expenses for both office spaces and rented equipment. These results are in line with the papers by Fukuda and Nakamura (2011 & 2013), which finds successfully recovering firms to be downsizing their workforce, which automatically means a lower need for rental expenses.

Fukuda & Nakamura (2013) found that restructuring a firm is very important during its recovery. From the results in this paper, I found this to be increasingly true when it comes to consolidations. The goodwill of recovering firms has increased from being the fiftieth most important distinguishing feature to being the second most important feature in the classification of recovering firms. This suggests the restructuring of recovering firms can be done via takeovers, especially during times with historically low interest rates. This is especially interesting considering existing literature found zombie firms to be a limiting factor for the expansion of healthy firms and the entry of new firms (Adalet McGowan et al., 2018) (Caballero et al., 2008). With the increasing importance of goodwill, this limitation can be expected to become even more relevant.

The results further relate to the findings of De Martiis et al. (2021). Their paper found common stock to be important for the recovery of zombie firms. The results below, however, indicate stock-related factors to be unimportant for a firm's recovery. This contradiction likely can be explained by the manner of data usage in both my own and De Martiis' paper. The results in this paper focus on recovering firms, the data therefore contains many recovering firms. The paper by De Martiis et al. (2021) focused on zombie firms and used recovered firms as a sub-question. Their data, logically, mostly contained zombie firms. This unbalanced manner of handling the data might cause their model to be only filtering the most extreme recovery cases, giving very different results as the correctly classified recovering firms are very different to the correctly classified recovering firms in this paper.

The results further show recovering firms benefit from shifting their focus towards intangible assets and operations. This has, to my knowledge, not been found by other research papers but could be a very cheap and quick method for zombie firms to recover to a healthier status. Refocusing on intangibles has the added advantage of higher tax write offs in the short terms.

Next section elaborates on gathering and preparing all necessary data, followed by an outline of the XGBoost models and optimization, after which the results are gathered and discussed in both a quantitative and qualitative discussion.

Data

To investigate recovering firms, data comparable to the papers by Acharya et al. (2022) and De Martiis et al. (2021) is used. The data has been gathered from both Compustat's global annual fundamentals and its North American counterpart. All data ranges from September 2004 until the same month in 2022. This period has been chosen because it entails two periods in which the economy had an upwards trajectory (2004-2007 and 2014-2019) and two periods with a low conjecture (2008-2012 and 2020-2022). Prolonging this range is possible but lowers the data density for many firms.

In line with previous research papers, firms operating in the financial services sector are excluded from the total dataset.

The global and North American datasets are combined into one dataset based on the 294 variables included in both the North American and global dataset. Non-overlapping variables are being dropped. Only firms with a valid identifier (*gvkey*) and reporting dates (*datadate*) are kept as the lack hereof results in many missing values in later steps.

This research focuses on Europe and North America, thus all observations from other countries are removed. Firms from 48 countries in Europe, plus firms headquartered in North America are kept. Europe is defined as member states of the Council of Europe plus Belarus and Russia. Belarus and Russia are included because they, especially Russia, have been important players in the European markets before declaring war on Ukraine. North America only entails the United States and Canada.

In some cases, firms report their fundamentals in more than one country, especially firms active in both North America and the Eurozone tend to report their financials in both Dollars and Euros. These firms are filtered based on the location of their headquarters. E.g., if a firm's headquarter is located in Belgium and reports its results in both Euros and Dollars, only its home country's currency is kept. Also, firms with headquarters located outside Europe and North America are dropped as important parts of these firms fall outside this paper's regional scope. In total these two measures deleted 1867 observations spread over a few hundred firms.

After this step, the dataset is checked for duplicates based on their identifiers and reporting dates, which resulted in zero duplicates. Several variables are deleted based on the number of missing values. Machine learning models are relatively good at coping with missing values but too many missing datapoints can cause biases. Therefore, any variable in which more than 65% of the datapoints are missing are not considered in the dataset. It is also important that the data holds enough information on various economic landscapes. Therefore, for a variable to be included in the final dataset, it must contain at most 50% missing observations in both North America and Europe. 102 different variables met these conditions and are included in the dataset.

These 102 variables are checked by hand to remove any unnecessary variables. Examples of this include variables with high correlation; intangible assets and tangible assets add up to total assets, including all three variables may distort the functioning of the model. In these cases, one of the variables is removed, in most cases the 'total' variable. 15 variables have been dropped because of this, resulting in 88 useful firm specifics.

Some firms are relatively young or go bankrupt relatively quickly. This results in a low number of observations for these firms. As these firms hold little information they can be removed from the dataset. The minimum number of observations that a firm needs in order to be included is set to four years of data. Four years is the cutoff since it is the timespan in which a firm can go from a healthy firm to a zombie firm to a recovered firm.

Definition of zombie and recovered firms

To predict whether a firm is healthy, a zombie, or even recovered, it is important to add a variable denoting this status. However, the definition of a zombie firm is not straightforward.

One of the more popular definitions is developed by Caballero, Hishi & Kashyap (2008). Their paper defines zombie firms as firms that receive subsidized lending, i.e. interest rates below the expected rates given a firm's borrowing quality. The main advantage of using this definition comes from the fact it can be used to properly investigate the effect of zombie firms on a sector or the general economy Definitions based on firm specifics make this research direction more difficult because an economic branch with a high number of zombie firms will always underperform if a zombie firm is defined by firm specific effects (Caballero, Hishi, Kashyap, 2008). Because this paper is not focused on spillovers coming from zombies, the main advantage of this first method is not very important. Further, banks and firms generally do not publicly state whether they are subsidized lenders or lending. This makes it difficult to observe the zombie status, especially considering a lot of public information needs to be added to the datasets spread over many different countries in a period with historically extremely low interest rates. The benefits of this method are therefore not worth the efforts needed to obtain all necessary information, which is why a second zombie definition is considered.

The second method was developed by Giannetti and Simonov (2013). Their paper benchmarked the interest rates of Japanese firms against the interest rates of firms with the highest possible financial rating. If a firm receives interest rates below the interest rates of the triple A firms, in combination with a below investment grade bond rating, as determined by the big three rating offices, it is considered a zombie firm. This method has the advantage of added ratings that prevent well performing firms from being classified as zombies because they incidentally receive very low interest rates. The disadvantage of this method stems from the fact firms can be misclassified due to accounting differences; a big write off or payment for a particular loan with low interest rates can cause a healthy firm to be classified as a zombie firm. Therefore, a third method was considered and finally chosen.

The third method uses the definition as developed by Banerjee & Hofmann (2018) and McGowan, Andrews, & Millot (2018). These papers define a zombie firm as a firm with an Interest Coverage Ratio (ICR) below 1 for two successive years in combination with a Tobin's Q below the median of the firm's industry over the same time path. This has been used as definition because the ICR shows a firm's capability to repay its debts while the Tobin's Q shows whether a firm's investments can be considered efficient or not. The main disadvantage of this definition is its lacking capability to determine spillover effects, especially in branches with a high presence of zombies. As this is not the focus of this research, in combination with the more firm specific approach, which is more in line with this particular research's goals, makes the definition from the two aforementioned papers the most suitable method for detecting zombie firms.

Tobin's Q is not included in Compustat's Fundamentals databases. Consequently, the 'Compustat daily price index' database is used to obtain each firm's number of outstanding shares (*cshoc*) and the closing price on each day that annual reports are released (*prccd*). The dates on which these two measures are obtained are preferably equal to the release dates of the annual fundamentals of each firm. These dates are generally on the same day each year for a specific firm. This causes some of these dates to be on Saturday, Sunday or on a holiday such as New Year's or Independence Day for US-based firms. Annual reports released on these dates caused missing values as stock exchanges are closed and no closing price can be determined. The missing trading days have been accounted for by shifting these observations' dates to the next trading day. Roughly 20,000 extra observations could be obtained this way. Cshoc and prccd are then used to calculate the market value of each firm each year, which can be used for Tobin's Q.

The ICR for each firm is calculated by dividing a firm's earnings before interest and taxes (EBIT) by the interest payments the firm must fulfil in the respective fiscal year.

By default, a firm is considered healthy if it is not a zombie. It is also possible to come back to life, creating a recovered firm. Once a firm is classified as a zombie, it remains a zombie firm until it meets the requirements to become a recovered firm. Recovered firms are classified in the opposite manner as zombie firms. To classify as 'recovered' a firm must have had a zombie status in the preceding year, in combination with an ICR above 1 and a Tobin's Q above the industry median in the preceding year.

Graph 1 shows the percentage of zombie firms and recovered firms in the dataset. It can be noted that there is a steep increase in zombie firms starting in 2008, can most likely be attributed to the financial crisis in this year. The percentage of recovered firms has a delay because one first must become a zombie before one can recover. In most years, some three percent of all firms recover from their zombie status while 9-12 percent of all firms are classified as zombies. These percentages suggest that the obtained dataset is in line with existing literature (De Martiis et al., 2018). The external validity is further confirmed by the US-only percentages (Appendix A) which are nearly identical to the percentages by the US government (FED, 2021).

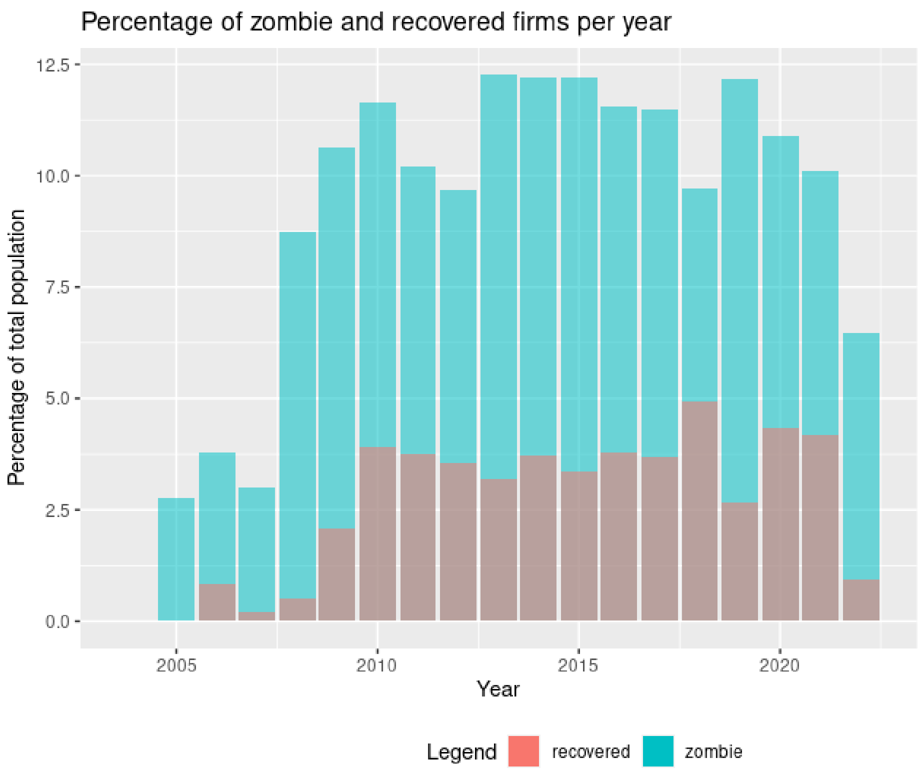

Graph 1: Percentage of zombie and recovered firms per year in both North America and Europe. Note that the bars are not stacked but show two different percentages. E.g., 2018 has 5% recovered firms and 9.8% zombie firms.

One of the biggest risks an XGBoost model can run into, besides outliers, comes from imbalanced data. Imbalanced data leads to models that seem to be very accurate at first glance but are useless when it comes to finding relevant features. This particular paper must account for imbalanced data because the vast majority of companies are either healthy or zombies. If the data is not balanced, XGBoost will most likely classify each data point as 'not recovered' because it will reach an accuracy at least 95% when doing this. To account for this, the XGBoost model uses data selections that consist of 50% recovered firms and 50% healthy and/or zombie firms. This will greatly improve the balance of the input data and will make the results more useful. Interestingly, no other paper can be found in this field that applies serious balancing to their datasets.

The risk of outliers in the final model has been mitigated by winsorization at 1%. Meaning for each numerical variable, the 1% most extreme observations are replaced by the 1^{st} and 99^{th} percentile. To improve the interpretation of the overall model and make to comparisons between variables easier for the model and its user, all numerical variables are scaled towards a normal $N(0,1)$ distribution.

After gathering, cleaning, and normalizing the dataset, the total dataset contains 79 explanatory variables and 1 dependent variable denoting the health status of each firm in each year. These 80 columns have 122,398 observations spread across 11,597 firms; 8,578 of which are from Europe and 3,019 from North America. A quick overview of this data can be found in table 1.

Year	# Firms in North America	# Firms in Europe
2004	984	3836
2005	180	3414
2006	275	4189
2007	1785	6135
2008	1739	6123
2009	1715	6890
2010	1800	7047
2011	306	5535
2012	1976	7822
2013	1936	7989
2014	1848	8176
2015	1712	8262
2016	225	6863
2017	1444	8234
2018	1300	8079
2019	1195	8043
2020	1147	7990
2021	1077	7713
2022	2	58

Table 1: The number of included firms in North America and Europe between 2004 and 2022.

Methodology

After cleaning and preparing the data, a machine learning model is used to classify each observation in the category ‘healthy/zombie’ or ‘recovered’. The use of such a model can be deemed important as it helps to discover patterns that are not immediately apparent to the human mind. Another factor lies in the size of the dataset. This paper will employ 80 variables, spread over more than 120,000 observations. This can be rather challenging for a traditional model, especially in the variable selecting part of the model. Using a machine learning model will help to overcome this problem.

XGBoost

Specifically, an XGBoost model will be used. A description of how this model works under the hood will be included in the next paragraph. XGBoost has three main advantages.

Firstly, XGBoost is known for its versatility; many hyperparameters can be set to make a workable model in many different situations, as will be elaborated on in the next paragraph.

The second reason to choose XGBoost rather than for example CART, which is its predecessor and more often used in existing literature, stems from XGBoost’s superior speed. XGBoost is mathematically optimized for both a computer’s CPU and RAM, whilst allowing for parallel computing; meaning XGBoost can run a specified number of trees simultaneously to employ one’s whole CPU rather than a single core at the time.

The third advantage XGBoost has over its competitors stems from its capability to work with missing data points which is important as some variables are not filled in for each company due to differing accounting standards in different nations, e.g., GAAP vs. IFRS. As explained in the data section, only a firm's health status and reporting date do not allow for missing values, the other 79 variables do allow for missing values which makes it important to choose a model that can manage missing data.

The advantages above are important when choosing a well performing machine learning technique for a given project, they are also amongst the reasons why XGBoost was used to win 17 out 29 challenges during Google’s Kaggle Competition, according to Kaggle.

The most important advantage XGBoost has over other models for this specific research lies in its capability to work with correlating data. XGBoost automatically removes perfectly correlating variables but also deals very well with partially correlating variables. It does this by focusing on specific variables for each iteration. If the model understands the effect of a variable, it will minimize the usage of that variable in future trees. Thus, correlating variables will be used independently from each other, ensuring little to no correlation effect in the feature importances. Further, the use of subsamples and colsamples, as will be elaborated on below, extends to this by not using all variables for each iteration, making it possible to analyze the behavior of certain variables even more independently. This is of great importance for this paper because balance sheet items will most likely be correlated; if a zombie firm has low profits, it will reasonably also have low investment expenditures and low asset changes.

Although XGBoost has been an industry standard over the last few years, it comes with a few disadvantages. XGBoost's main drawbacks lie in its sensitivity to outliers and the risk of overfitting (Choudhury, 2021). The risk of outliers means the model is relatively sensitive if any extreme values are included in the model, especially during the first few generations of the learning process. To mitigate this risk, the 1% most extreme values are winsored and the learning rate is set to conservative percentages.

The risk of overfitting stems from the fact that a model that has been trained on the same dataset with wrong hyperparameters becomes extremely fitted for its current purpose, generating high accuracies but low external validity. External validity is extremely important in an academic context and is ensured by scoring each generation of the model using a test dataset. The data will be subdivided into two sets; one includes 75% of all observations, all randomly selected, and the other contains the remaining 25% of observations. The first set can be used to train the models and to learn all the important features in the data, which will then be applied to the second set to test the actual usefulness and external validity. All results in this paper will be coming from this test set as it is the most reliable in the real world. Another method to mitigate overfitting is to set the hyperparameters conservatively to ensure the model does not have the freedom to build too many very small or irrelevant branches and leaves.

The last main disadvantage is not idiosyncratic to XGBoost but pesters most machine learning models. These models are so-called ‘Black Box’ models, meaning we know how it works, we know the input and the output, but the model is too vast to know why exactly the models are making certain classifications. This causes Black Box models to be very hard to interpret. To overcome this problem, newer methods have been developed over the years, including SHAP and several built-in functions in R. These methods analyze the steps taken by XGBoost and can output, amongst others, the most important factors in the decision-making process of the models, making the output much more useful.

The disadvantages of XGBoost can be solved with relative ease for this research. Because of this, in combination with the many advantages XGBoost has over other machine learning techniques and traditional methods, XGBoost is considered the most suitable fit for this paper’s data and research.

A brief explanation of the XGBoost model

In this part, XGBoost will be briefly explained. Firstly, an explanation with its intuition is discussed, followed by a bit more in-depth explanation.

XGBoost, in its fundaments, starts as a random forest model. Random forests build multiple decision trees based on input data. When an object needs to be classified, the trees all have 1 vote and will vote based on the classification resulting from their tree. The observation will then be classified in the class that received most votes. XGBoost, as the B in its name suggests, Boosts the next generation of trees. Boosting means the model gives more weight to observations which were classified in the wrong class and less weight to observations that are easy to classify. Boosted models therefore quickly learn the most important facts about a dataset and do not only account for the more basic data features but also for patterns that are harder to discover. On top of regular boosting, XGBoost applies Gradient Boosting. Gradient boosting applies a gradient descent procedure based on the low of the last iteration. Then, it focuses on the decisions that caused the most loss by adding them one by one to each new tree, this to minimize loss. XGBoost goes one step further than a regular Gradient Booster with the addition of many different hyperparameters and settings that can be used to optimize the model for virtually any data related problem. These hyperparameters will be discussed in the next section.

The paragraph above is an intuitive but very broad summary of what XGBoost entails. In the following part, a deeper dive will be taken. The XGBoost model, fully known as the 'Extreme Gradient Boosting' model is developed to improve on the existing 'Gradient Boosting' model. Both these models are boosted versions, i.e. upgraded versions of the original Random Forest techniques developed by Tin Kam Ho (1995)[1]. Tin Kam Ho's introduction to Random Forest models caused the world of algorithms and decision models to take big leaps over the past decades. The modern XGBoost goes even further and works as follows:

First off, the XGBoost classification model used in this paper, as most models, is used for supervised learning problems. Meaning it uses training data with multiple variables (X_i), with the goal of predicting the dependent or, more formally, the target variable (Y_i). The models in this paper use a Y_i consisting of the categories 'healthy&zombie' (denoted by zero) and 'recovered' (one).

To discover the best hyperparameters (θ) for each year, regularizations terms (Ω) are added to cap the model's complexity and to limit the chance of overfitting. Said regularization terms will be discussed below. The regularization terms are accompanied by a loss function (L). The loss function shows how much information is lost in the prediction with respect to the training data. The loss function for classifications is traditionally chosen to be the 'multiclass squared error' (m-error). This paper follows this rule of thumb and thus uses the loss function given by the original XGBoost documentations (Chen et al., 2016):

$$L(\theta) = \Sigma\left(Y_i - \hat{Y}\right)^2 \qquad (1)$$

The Ω and loss term are combined into one objective function and minimized. This minimization is important because a small Ω ensures the trees are not too complex and thus overfitted/biased. A small loss function ensures more important information from the training set is included in the model, making the predictive power higher. The principle of an objective

[1] The first fundamentals of Random Forests were developed by Babylonian and/or Egyptian mathematicians in c. 1500-2500 B.C. (Jean-Luc, 2012).

function consisting of a combination of the Ω and loss function is considered a solution to the bias-variance tradeoff commonly known in the data science environment.

With this basis, the second step of building an XGBoost model is to use decision trees. Each observation is classified, using these decision trees, into their own group, called a leaf. Each leaf of each tree is then evaluated using a prediction score. These scores are combined to calculate the overall scores of a specific tree using an objective function which is updated by a constant and so-called Taylor expansion. Taylor Expansion is a mathematical process used to update and expand the existing object function to include relevant information from the current tree so it can be used when building the next tree. The exact workings of this process can be found in the original paper of the XGBoost model (Chen et al., 2016).

In the third step of building an XGBoost model, the regularization terms are added to the objective function using a mathematical process which outputs a complexity score. This step is followed by several other statistical necessities which are beyond the scope of this paper but can be found in the original paper by Chen et al. (2016). The last step considers all statistics, and simply adds them together in a 'structure score' which shows a tree's performance based on its complexity and performance given the structure of the tree.

To find the best tree possible, all different possible parameter combinations should be used, which would result in an infinite number of possibilities. Hence, the addition of each split is evaluated using the 'Gain' , given by:

$$Gain = \frac{1}{2}\left[\frac{G_L^2}{H_L+\lambda} + \frac{G_R^2}{H_R+\lambda} - \frac{(G_L+G_R)^2}{H_L+H_R+\lambda}\right] - \gamma \qquad (2)$$

This formula takes the prediction scores of both the left and right leaf created by an additional split (G_L and G_R) divided by their loss functions (H_L and H_R). These terms are added together and. Then, the scores of the original leaves (third part of formula) in combination with a regularization term (λ), which is set to make sure trees do not grow too deep are subtracted. Lastly, the γ is subtracted. The gamma is included as a minimum required improvement that must be gained from an additional split in the tree. This is a way of keeping the trees smaller and more workable, and is generally referred to as 'pruning'. If the Gain equals a score higher than zero, it is considered a valuable addition to the existing tree.

Setting hyperparameters

XGBoost has many parameters that can be included as regularization measures to make a better tree. To optimize this process, this paper employs a grid search and builds a new model for each combination specified in the parameters grid. To get a better understanding of this process, the next part will briefly describe each hyperparameter and their used values. First off, however, the initial settings of each tree will be briefly discussed.

Classification and classes

The XGBoost model will be classifying each observation using the multiclass method. This means it will determine each observation's probability of being either a recovered firm or a zombie/healthy firm. It will then classify the observation into the class with the highest probability.

M-error

While training the trees, the 'multiclass error rate', better known as m-error, will be used to evaluate the performance of each individual tree. The m-error is simply calculated by dividing the number of wrongly classified observations by the total number of classified observations. A lower m-error means the model is misclassifying less observations and is thus improving. Each tree outputs two m-errors; one based on the training set and another based on the test set. This paper uses the m-error from the test set to evaluate the trees as it helps to find an optimal tree that is more externally valid.

Time varying variables

It is possible for variables to be important during certain periods in time. When answering the research question, it is necessary to investigate the key factors of a recovery over time. Also, this will make the overall results more robust against features with high peaks in importance during specific periods, which might skew the overall classification processes.

To do this, the models are run from 2009 onwards, using the current year plus the two upcoming years. E.g., the first model is built for the years 2009-2011, followed by 2010-2012, etc. Combining three years for each model is important to have enough observations per year to work with, while retaining the advantages of running shorter term models.

The dataset contains data starting in 2004, the first five years of data however contain many missing observations. These turned out to be too small for a reliable XGBoost model in which cross-validation and subsampling can be used for optimization, as will be further explained below.

In total, 11 periodic models are built and optimized using a grid search. The grid search uses the hyperparameters explained in the next paragraph to get the best results for each model.

Setting hyperparameters for Grid search

<u>Similarity score</u>

The similarity score uses the residuals of each leaf, controlled for the number of observations, to see whether the addition of an extra split results in two very similar leaves or not. If the leaves are very similar, it is not very useful to include them in the trees. The similarity scores need to have a starting point to work with, which is set at 1 divided by the number of classes, 0.5 in this paper's case. To determine the minimum required difference between two new leaves, the lambda is set as a regularization term.

<u>Lambda</u>

The lambda helps to lower similarity scores, which is especially important in leaves with few observations as these can easily be influenced by outliers in the data. A higher Lambda results in a lower chance of overfitting due to outliers and vice versa. As outliers are removed from the data, lambda does not have to be set too conservatively. The default lambda equals 1, but most literature uses values between 0.4 and 3.2. This paper will include a lambda range of [0.4, 0.8, 1.6].

<u>Minimum child weight</u>

Another parameter limiting the effect of overfitting is the 'minimum child weight'. This hyperparameter sets the minimum number of observations a leaf must include for it to be considered. If a tree contains many leaves with very few observations, it risks lacking external validity as micro decisions are made based on -likely- irrelevant information within the training set. It is therefore important to prune the tree by adding the minimum child weights. The weights in this paper are as follows: [1, 5, 10].

Gamma

Another pruning measure can be found in the gamma, which is subtracted from the Gain. If the gamma is not included, the trees would grow almost indefinitely. A low gamma logically results in very deep trees as there is a low bar for the splits to be included. This paper includes a relatively high number of observations and variables. Therefore, gamma is set at [0.8. 1.6, 3.2] so it allows for relatively deep trees.

Learning rate

Even - or especially - computer models can have incorrect conclusions. Therefore, only a percentage of each iteration's findings must be included in the next tree. If all information were concluded, the model would risk finding a local optimum. The learning rate of the models, also known as ϵ, will come from the following grid: [0.01, 0.1, 02]. The default learning rate is 0.3 but given the vast dataset, the model must learn slower for it to accommodate all different observations.

Cross Validation

To make a more externally valid model, each combination of parameters is cross validated. This will be done in 2-fold, meaning the model draws two separate training and test sets for each individual tree. The model trains and tests on both these sets using the current combination of hyperparameters. This helps the output of the model to be more robust against outliers and other irregularities in the dataset. To do this properly, a subsample rate and a colsample rate are set.

Subsample and Colsample

The subsample rate in XGBoost denotes the percentage of observations used in each training set for the cross-validation process. This paper uses a range of [0.5, 0.75], meaning either 50 percent or 75 percent of all rows in the training set will be used. The Colsample rate parameter uses the same range but for columns, each column has a 50 or 75 percent chance of being used. This helps XGBoost to cross-validate but also limits the risks of having distorting variables included in the final model. On top of that, partially correlated variables will function independently from each other because the model will regularly draw only one of the correlating variables. This is a tremendous help for the importance analysis and is considered to be one of the reasons XGBoost is capable of working with correlating variables very well.

Max Depth

A very important hyperparameter is the Max Depth, it determines how many layers each tree is allowed to be. Setting the Max Depth very low results in shallow trees with few leaves. This has the advantage of being easily interpretable but also creates the risk of missing important splits and may result in a very general model that is less accurate. Setting the depth very high creates a deeper tree with many splits, which can lead to overfitting as it can include splits with very little or even wrong information. This paper uses a relatively broad range for its maximum depths, which is [4, 7, 10, 15].

Number of trees and early stopping

XGBoost will not stop learning and building new trees until it is told to do so. In case of building too many trees, XGBoost will use each possible bit of information to reach exceedingly small gains in accuracy. This consumes a lot of time and will eventually run into the risk of using information that is only relevant within the specified dataset, i.e., overfitting.

It is possible to find an optimal tree, after which the model declines in performance if more trees are built. This is accounted for by adding an 'early stopping parameter' to the model. The early stopping parameter checks if there are any significant gains in performance every 50 trees. If the performances of the trees started to decline in these 50 rounds, the best performing tree will be saved as the final model for a given parameter set. Many projects use 10 or 50 as the value for this parameter. This paper uses 50 because setting low values may cause the model to find a local optimum while 50 trees can account for these temporal dips in performance.

Because relatively low learning rates are used, the number of trees must accommodate for this by being relatively high in order to give the model a chance to reach an optimum in time. Therefore, each model will train until it reaches its optimum or until it has built 1000 generations of trees. 1000 has been empirically chosen, adding more trees resulted in very little gains, suggesting that the model is running into overfitting problems, as discussed earlier.

Grid Search

As discussed before, this paper employs a grid search to find the best possible hyperparameter combinations. It does this by running each model for each possible set of hyperparameters. Given the hyperparameter above, the grid search will go through 1296 different sets of hyperparameters for each period. The results of each iteration are saved, and the best performing hyperparameter set is used to build the final model for each period. An overview of the sets used for the final models can be found in table 2.

Table 2 shows the optimal models prefer a low learning rate with relatively deep trees and low minimum child weights. In combination with the low gammas and lambdas, the trees produced by these models can be expected to be fairly large.

Model	***Subsample***	***Colsample***	***Max Depth***	***Eta***	***M.C.W.***	***Gamma***	***Lambda***
2009-2011	0.75	0.75	15	0.1	5	0.8	0.8
2010-2012	0.75	0.5	10	0.01	1	0.8	0.4
2011-2013	0.75	0.5	10	0.01	1	0.8	0.8
2012-2014	0.75	0.5	10	0.01	1	0.8	0.4
2013-2015	0.5	0.5	7	0.01	1	0.8	0.4
2014-2016	0.75	0.5	10	0.01	5	1.6	0.8
2015-2017	0.75	0.75	10	0.01	1	0.8	1.6
2016-2018	0.75	0.75	15	0.01	1	0.8	0.8
2017-2019	0.75	0.5	15	0.01	1	0.8	1.6
2018-2020	0.75	0.75	15	0.01	5	0.8	0.8
2019-2021	0.75	0.75	10	0.01	5	0.8	1.6

Table 2: The hyperparameters found to be optimal for each period's XGBoost model. All trees have a maximum of 1000 iterations. M.C.W. stands for Minimum Child Weight.

With these optima, it is possible to derive the results necessary for answering the research question.

Results

After building all 14,296,000 trees, the best tree for each period has been found, as shown in Table 3. Appendix B contains one of the trees as an example, making it perfectly clear why XGBoost dearth interpretability. The kappa for each period is moderately high, indicating there is a low chance that the predictions are determined by chance. The significance of each model is also included in the table to evaluate whether the two classes are different. The models are all significant with a p-value < 0.0001.

The recovery prediction rate for each period denotes the percentage of recovered firms labeled as recovered firms by the XGBoost model. This percentage is consistently over 80%, showing the model can define and predict recovering firms rather well. There is a minor increase in accuracy over time which is most likely due to an increase in data availability in recent years. The overall recovery prediction rate is roughly five percent lower than the zombie prediction rate by De Martiis et al (2018) which used a more traditional CRANE model to discover zombies. The difference in accuracy can possibly be explained by the higher availability of zombie firms; there are about four times more zombies than recovered firms each year. Another explanation for their higher prediction rates might lie in the balance of the dataset. The data in this paper is balanced into sets with 50% recovered and 50% non-recovered firms. Initial runs with unbalanced data resulted in much higher accuracies. The problem with those initial models, however, lay in the recovery prediction rates, which were much lower. This turned out to be caused by the low percentage of recovered firms which made it tremendously difficult for the XGBoost model to find splits that work better than simply assigning most observations to the non-recovery class. Adding a more balanced data sheet decreased overall accuracy but greatly improved the prediction of smaller classes, whether it be zombies or recovered firms.

The recovery prediction rate of course only gives a broad view of the results with low interpretability, which is why the most important features for each period are analyzed for a more in-depth review.

Model	Kappa	Recovery prediction rate	significance
2009-2011	0.5437	0.8053	****
2010-2012	0.6438	0.8262	****
2011-2013	0.6614	0.8651	****
2012-2014	0.6473	0.8275	****
2013-2015	0.6227	0.8217	****
2014-2016	0.6587	0.8616	****
2015-2017	0.6741	0.8472	****
2016-2018	0.6928	0.8541	****
2017-2019	0.6677	0.8513	****
2018-2020	0.7051	0.8606	****
2019-2021	0.6575	0.8354	****

*Table 3: Initial performance results in each period. Significance levels: * at <0.1, ** at <0.05, *** at <0.01, **** at <0.0001*

Graph 2 visualizes the most prominent features for each period. The highest possible rank for a feature is 79 with a lower bound of 1. This graph only shows variables which have been in the top 3 of most important features at any given point in time. Table E1 in Appendix E shows the top 20 most important features measured as gain for each year. Graph 2 visualizes some of the most interesting patterns in the data. One example of this can be found in the retained earnings (re). Retained earnings is the most important feature in almost every year, which makes sense given that more retained earnings can be expected to be associated with financial wellbeing. The retained earnings of recovered firms are similar to those of healthy firms with means of 11.75 million and 12.78 million, respectively. This is in sharp contrast to the mean retained earnings of zombie firms, which is 1.18 million. The relation between retained earnings and a firm's recovery becomes also clear when looking at their Pearson's product-moment correlation; the retained earnings of recovered firms are negatively correlated with the retained earnings of zombies at a p-value of 0.00001, whilst no significant correlation can be found between the retained earnings of healthy and zombie firms.

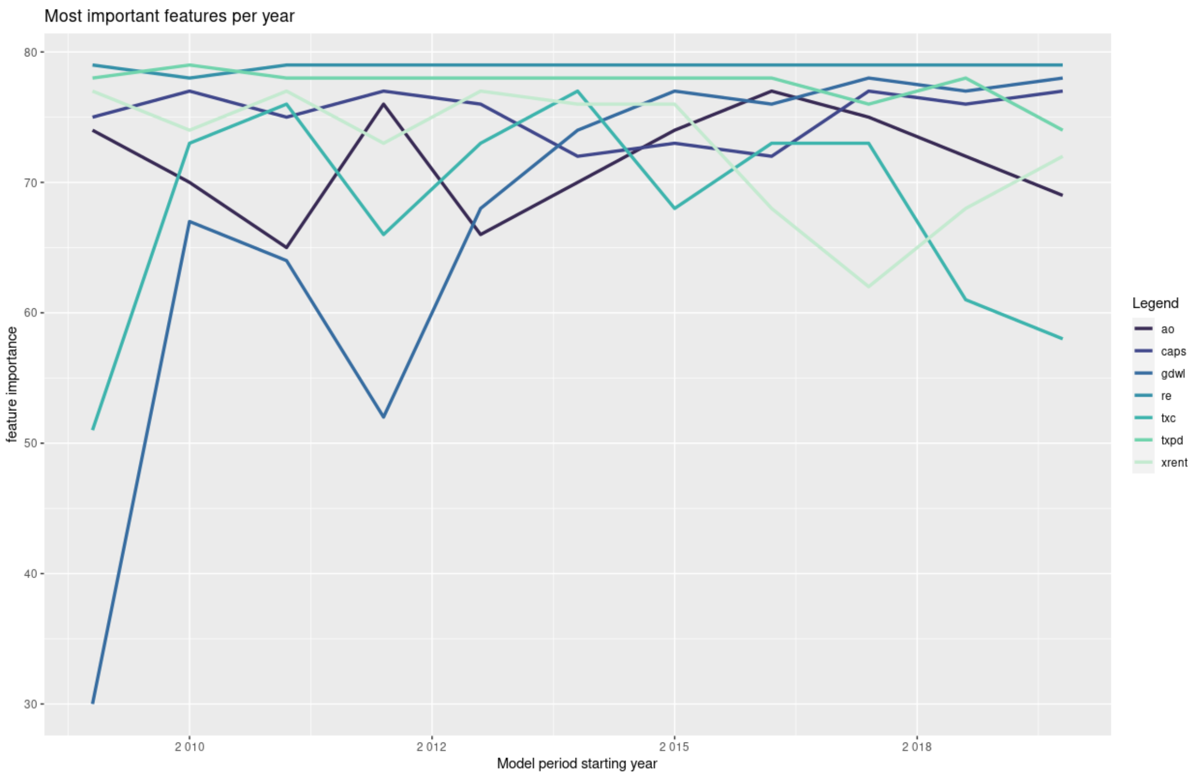

Graph 2: Most notable features for each year. ao; Assets other, caps; Capital Surplus, gdwl: Goodwill, Re: Retained Earning, txc: Income Tax Current, txpd: Income Tax Paid, xrent: Rental Expenses

Another interesting pattern in the important features can be found in goodwill (gdwl). The importance of a firm's goodwill grew from the 50th most important feature in the period 2009-2011 to the second most important feature in 2019-2021. In this period, its importance, measured as gain, tripled; almost 3% of the later models' performances can be attributed to goodwill. Goodwill almost exclusively comes from consolidations. The increasing importance of goodwill indicates zombie firms that recover increasingly manage to do so via mergers and acquisitions. At first glance this seems unlikely as zombie firms normally lack sufficient funds to properly invest in their own firm, let alone other firms. To further investigate this, Shap values are calculated for the models, one of which is shown in graph 3.

Shap values indicate if certain values have a positive or negative effect on the predictive performance of the XGBoost model for a specific class. Graph 3 shows for goodwill, higher values in the normalized data are associated with a negative Shap value. This indicates that firms with high goodwill values have a negative effect on the predictive performance of XGBoost for the recovery class. Thus, recovering firms are often characterized by lower goodwill values. It is important to note this does not mean that zombie firms with higher goodwill values have a greater or smaller chance of recovery, the Shap values merely indicate that firms with higher goodwill values make it more difficult to classify a firm as recovered because it is not considered a standard characteristic for these firms.

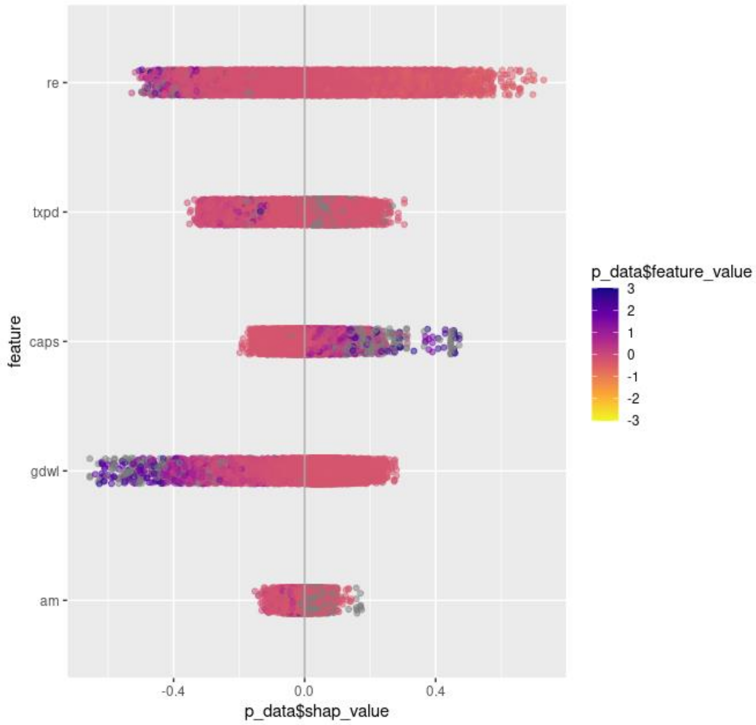

Graph 3: Relation of Shap values and feature values between 2018 and 2020 for the top 5 key features. Shap plots for other periods show the same patterns. Feature values are normalized between –3 and 3.

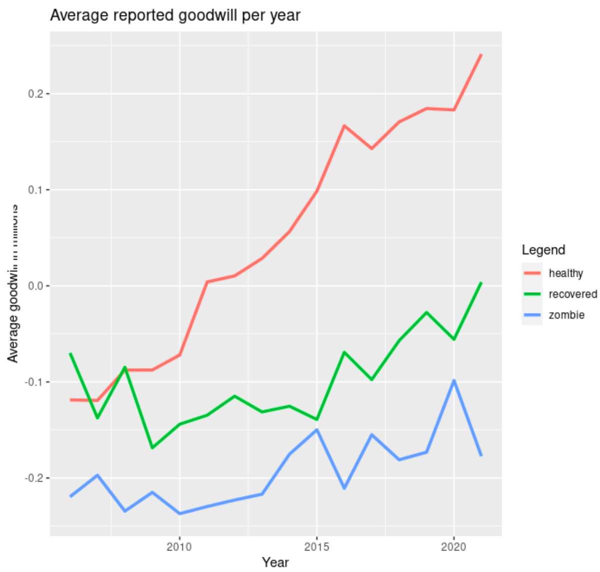

Graph 4: Average Goodwill per year. Goodwill data is normalized to N(0,1) where 0 is the mean value of Goodwill between 2004 and 2022

Graph 4 shows the average goodwill of recovering firms is higher than their zombie counterparts, but it remains much lower than the goodwill values of healthy firms in most years. This can possibly be explained by the historically low interest rates during the observational period. Borrowing money became a lot cheaper, making it possible for firms to indebt themselves to acquire another firm without much trouble.

On top of the historically low interest rates for the whole economy, zombie firms also benefit from credit misallocation from zombie lenders, making it possible for zombie firms to borrow money against interest rates below the efficient market rates (Caballero et al., 2008). Healthy firms still outperform recovering firms, which can be explained by their greater number of liquid means and the fact non-zombie lenders are most likely more willing to lend these firms the required funds for a takeover. However, monetary policies in the capital structure channel by the European Central Bank can be a reason for non-zombie lenders to seek more risk (Grosse-Rueschkamp, Steffen and Streitz, 2019). This has likely been playing in the hands of zombie firms in recent years.

Goodwill's importance, shap values and the average trend over time found in this paper suggest that an increase in goodwill, i.e., taking over other firms, is a distinguishing characteristic of recovering firms and it might be beneficial in some cases, as indicated by the higher values compared to zombie firms. However, the results cannot be used to make a convincing case for takeovers to be a good method for overcoming a firm's zombie status.

Another notable data feature is the amortization. The amortization is not shown in graph 2 but can be found in Appendix D. The amortization's gains grew by 0.45 percentage points, making it the fifth most important feature in the last period with a gain of 2.29%. The importance of amortization suggests that recovering firms use their intangible assets to a greater extent than their zombie rivals.

The top 3 most important features further indicate some outgoing tax indicators show signs of declining importance. The outgoing tax payment cash flows (txpd) are traditionally the second most important data feature but has been declining in recent years, while the current taxes (txc) also suffer from a downward trend in their gains. The average current and paid taxes per year are shown in Graphs C1 and C2 in appendix C. These graphs make the importance of the taxes paid in the decision-making process of the XGBoost models very clear.

Recovering firms pay on average less taxes than healthy firms but their trend over time is much closer to that of healthy firms as compared to zombie firms.

The rationale behind the decreasing importance of taxes does not become immediately clear from these results. A –speculative- explanation could be the introduction of several tax measures during the covid outbreaks in the later years of the data period. These tax reducing measures may have compounded on the credit misallocation which often is cited as one of the main problems causing zombie firms to survive. Governments generally gave tax breaks to all firms without considering their financial wellbeing, giving an opportunity for some firms to stay afloat even though they would have gone bankrupt in 'normal' circumstances. These tax measures may therefore have been a less important factor in the success or failure of a firm, creating less importance in the XGBoost models.

Also, graph C1 shows that during the crisis years of 2007-2009 the taxes paid for recovering firms fell much harder than taxes paid by healthy firms and became nearly identical to the taxes of zombie firms, this is also indicated by the lower current taxes in 2009 shown by graph C2. This is of course an indicator of their profitability during these years and might suggest that during crisis years recovering firms are tax-wise and thus profit-wise much closer to zombie firms. A closer relationship with zombie firms during crisis years suggests that in order to recover a firm during 'good' years, maximizing profits is a better strategy than it would be for a firm that wants to recover during more uncertain times. Please note that this explanation is purely based on only two periods with merely visual indicators, it should therefore be taken with a grain of salt. It would, however, be interesting for future research to do an event study regarding the relationship between taxes, profitability and recovery rates during crisis years.

The last interesting aspect of the most important features lies in the capital surplus (caps). To recover as a zombie firm, increasing your capital surplus is among the first things that come to mind as it is one of the reasons companies end up as zombies in the first place. Confirmation of this can be found in the importance of capital surpluses in the results. Over the whole time period, capital surpluses are in the top 10 most important features. This phenomenon is further endorsed by the average surpluses per firm type over the years (Appendix F) which shows the capital surplus of a recovering firm to be very close, and often even higher than that of healthier funds, while zombie firms are struggling to keep their funds up. Both Pearson's product-moment correlation and Spearman's rank correlation show that the capital surpluses are correlated with a significance level of 1% and coefficients of 0.943 and 0.932 respectively. The correlation between the capital surpluses of recovered and zombie firms results in coefficients which are roughly 40% lower and significant at 'only' 5%.

Graph 5 shows the least important characteristics of all firms classified as recovered by XGBoost for each period. Each variable included has been in the bottom 5 important features for at least one year. Unsurprisingly, both redeemable and unredeemable stocks (pstkn and pstkr) are not very important indicators of a firm's recovery. Excise taxes also do not play a big role in a firm's recovery, which makes sense as it is not part of the current firm but is purely collected and transferred for legal reasons.

There are, however, more interesting results in this graph; namely the turnover (sale) of a firm. The turnover of a firm does not hold any significant meaning when trying to classify recovered firms as recovered firms. This suggests the turnover of a recovering firm is not contributing much to the upwards trend of the firm; or at least it is not something that distinguishes recovering firms from zombie and healthy firms. The lacking importance of sales in combination with the importance of retained earnings further suggests that firms that want to recover should focus on the maximization of their profits rather than maximizing the size of their firms through turnover increases.

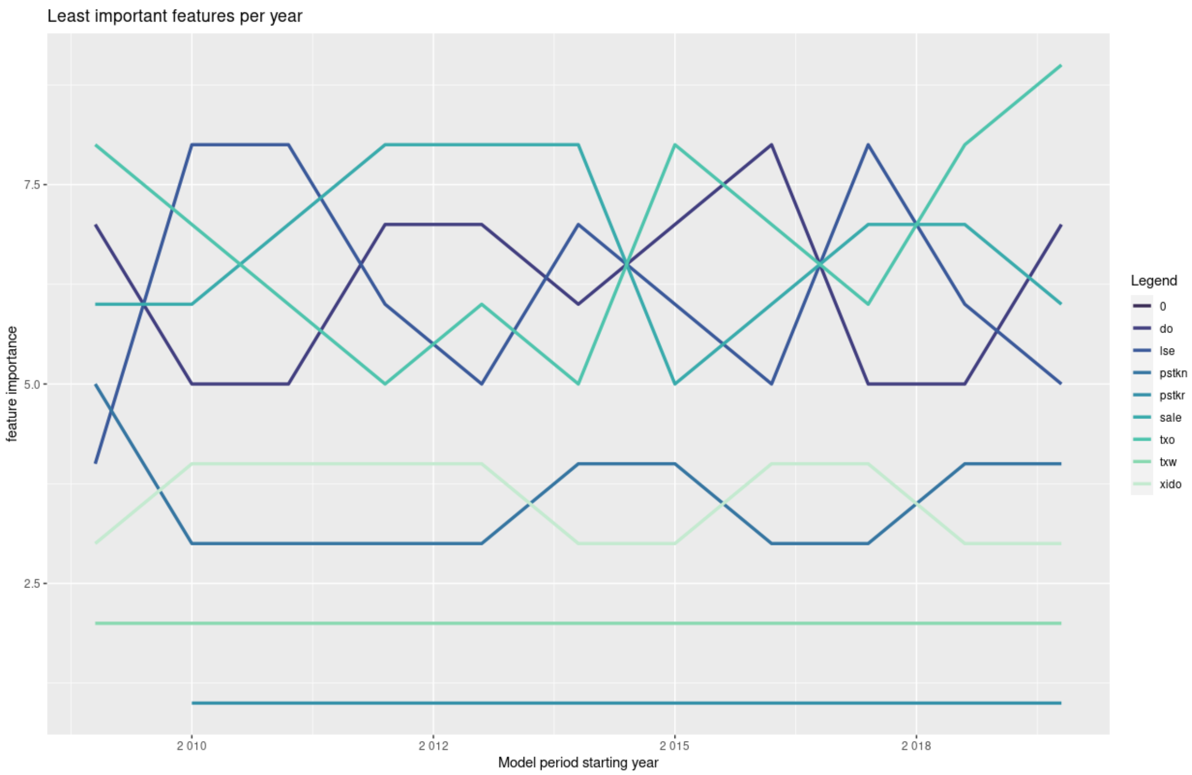

Graph 5: Least important features for each year. The importance is shown as rank, a feature importance of 1 means it is the number 1 unimportant feature for a given year. Do: discontinued operations, lse: Total Liabilities and Stockholder Equity, pstkn: Preferred Stock Nonredeemable, pstkr: Preferred Stock Redeemable, sale: Sales/Turnover, txo: Other Income Taxes, txw: Excise Taxes, xido: Extraordinary Items from Discontinued Operations.

Another aspect of the least important variables can be found in the discontinued operations (DO) and the extraordinary items rising from said discontinued operations (XIDO). These two balance sheet items together contain all profits and losses arising from closing or selling a part of a firm. The lack of distinguishing power is visualized in graph 6 which plots the average profit (loss) per year coming from discontinued operations. The trends in this graph are nearly identical for each category. The only difference can be found in the lower volatility of healthy firms, which can be explained by the fact healthy firms have more observations and are thus less prone to moving averages resulting from a few big or small observations.

The lack of importance for DO and XIDO is rather interesting as selling a part of one's operations can be used to generate capital relatively quickly whilst also forcing a firm to better focus on the remaining company. The lower importance can be explained by the occurrence of credit misallocation; the operations a zombie firm wants to sell or discontinue can be worth more to the zombie firm than to other firms because the zombie firm has the benefit of zombie lenders granting them low interest rates, making the funding of the firm relatively cheap. The zombie firms evaluate the operations they might want to discontinue based on lower financing costs, which creates a situation where the operations are economically viable with positive net present values for the zombie firm, although these operations should not be viable in a non-zombie economic paradigm (Caballero et al., 2008). If a healthier firm wants to buy a part of a zombie firm, it faces higher financing costs, making its valuation lower. This creates a discrepancy in the valuations made by healthy and zombie firms, making negotiations difficult and acquisitions by healthy firms less likely.

DO and XIDO do not only arise when selling parts of a firm but also when simply bringing parts of an organization's operations to a halt. This might be a strategy to free up cash flows for more vital parts of a company. The low importance can be explained by the same reasoning as presented for the takeovers above: Zombie firms have low financing costs due to credit misallocation by zombie lenders, creating net present values that are positive although these operations should be considered economically inefficient.

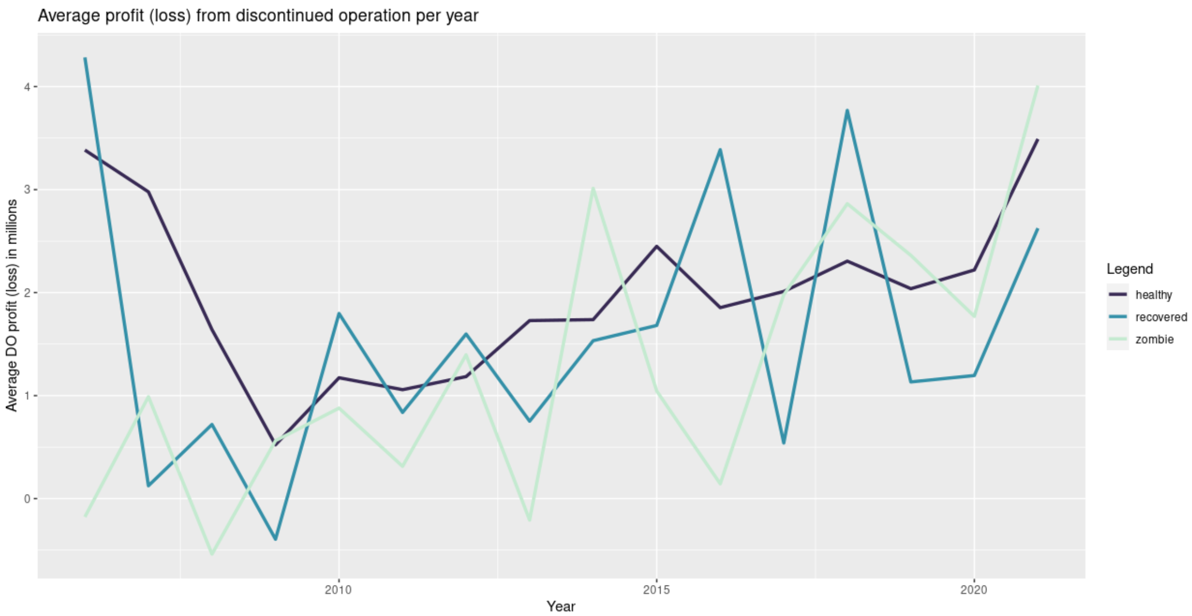

Graph 6: Average profit or loss from discontinued operations per year.

Qualitative measures

The qualitative measures that can be taken by managers following the results are of great importance when answering the research question. The results above suggest that a manager in a zombie firm should not focus on discontinuing operations or maximizing sales rather than maximizing profit. Zombie managers should, given these results, focus on increasing retained earnings, amortization and capital surpluses whilst lowering rental expenses and taxes. These five measures will be discussed in the next paragraphs.

Increasing retained earnings

As discussed before, the retained earnings are the main deciding factor for recovering firms, suggesting recovering firms can optimize their retained earnings in manners that other zombie firms cannot. To increase a firm's retained earnings, one must either increase net income via revenue, increase net income by cutting costs, or decrease shareholder dividends. Increasing the net income of a zombie firm via an increase in revenue comes with two main problems. Firstly, it most likely cannot be done using investments as zombie firms by definition do not have enough capital to invest and even if they can invest, increasing sales hardly contributes to a firm's recovery, as shown in graph 5, so the firm should only invest in profit margin increasing measures. The second problem lies in the increase of a zombie firm's profit margin, zombie firms cannot simply increase their prices because their ability to survive with low prices is cited as one of the reasons a zombie firm can persist in the market (Nakamura, 2016). Thus, to increase profit margins, a zombie firm must focus on increasing the firm's productivity and innovations. To do this, the firm must invest in new equipment or other capital goods, which brings us back to the first problem.

It is thus unlikely for a zombie firm to recover using retained earnings if it only focuses on increasing net incomes via investments and revenue. Banerjee et al. (2020) found that zombie firms hardly pay dividends, ruling out a decrease in dividends as a possible measure to recover one's firm. What remains is the option to decrease costs. Decreasing costs can be done in various ways and heavily depends on the type and structure of a firm. Most firms have labor wages as a huge cost factor. Saving money on wages, however, is likely to lead to lower productivity and therefore lower profits, as is argued by Yelen (1984). Since zombie firms tend to have low productivity numbers, lowering wages is not a good measure to cut costs and increase the retained earnings. In existing literature, restructuring a firm's employees is found to be a better measure to decrease costs. For example, Fukuda & Nakamura (2013) found reducing the number of employees in a firm to be one of the key factors for Japanese firms that went through a zombie faze. They further find that removing positive incentives for the remaining managers in a firm has a negative effect on the chances of revival.

Thus, to increase the most important factor in a zombie's recovery process, being retained earnings, the managers of the zombie firm do not have to focus on increasing sales, investments or cutting dividends, but should focus on cutting costs. Effectively cutting costs can be done by shrinking the workforce while not cutting the wages and positive incentives for the remaining employees.

Increasing amortization

Although amortization has not been found to be one of the top drivers in recovering firms, its importance is steadily increasing over time. It could therefore be used by zombie managers who want to optimize their recovery plan. Amortization cannot simply be increased by investments because zombie firms are not capable of financing such investments. It is therefore likely that the higher amortization levels of recovering firms stem from the type of firms. This would mean that firms with a high focus on tangible assets, such as manufacturing plants, have more difficulty recovering than those with a more intangible focus, such as a design office.

For many firms it might not be possible to shift their focus to a completely intangible structure, but the available intangible assets can be an opportunity for firms to capitalize on parts of their firms that can be used without many necessary recourses and funds.

Increasing capital surpluses and lowering rental expenses

The increase in capital surpluses is one of the biggest issues for a zombie firm as it bars them from investing in their company and limits the ability to pay off debts. Increasing a firm's capital surplus can be done by decreasing the capital outflow, which can be done by cutting costs, as discussed in the paragraph about increasing retained earnings. A zombie firm can also increase its capital inflow to become a recovered firm. Fukuda & Nakamura (2013) found a method for this. Their paper found selling fixed assets to be one of the main contributors in the recovery of a zombie firm.

Increasing capital surpluses, however, should not be done by discontinuing operations as the XGBoost models found discontinuing parts of a firm to be one of the least important factors in a firm's recovery.

The importance of capital surpluses might also be tied to the importance of rental expenses. If a firm sells its fixed assets and downsizes the company, lowering the number of rented buildings and assets is a logical next step for a firm. Decreasing rental expenses will have a positive impact on a firm's capital surplus. Decreasing rental expenses has become much easier in recent years because the recent pandemic has normalized working from home, meaning it is possible for distressed firms to cut in the number of offices and other physical matters that can be -partially- done from home.

Therefore, managers of zombie firms can increase their capital surplus by selling fixed assets, downsizing the company and by making use of the recent increase in innovations surrounding working from home to cut rental expenses.

Lowering taxes

The taxes that must be paid by zombie firms are a major burden on a firm's recovery. Lowering taxes cannot be done easily as managers quickly stumble on both legal and ethical boundaries. Besides, the international character of this research means there are many different taxation guidelines, making it difficult to analyze the best method to lower taxes for a firm that wants to recover from its zombie status. In general, many companies keep their taxes low by using international treaties, setting up systems in tax heavens, increasing depreciations for short term gains and many other bookkeeping tricks.

The measures taken to benefit from taxes should be looked at on a firm-by-firm basis and could be a very interesting topic for further research. The results of this paper can be a good starting point for such research because an increased amortization, downsizing the firm’s workforce and selling fixed assets can all be short term tax write offs that help the zombie firms to reach a recovering status.

Conclusion

Existing literature has established an elaborate picture of what the main characteristics of zombie firms are. Many different methods have been employed for this, including some standard machine learning techniques to investigate what causes zombie firms to reach their zombie status over time. Hardly any research used more modern and versatile machine learning techniques to capture the drivers of a firm's recovery. To find what managers in zombie firms can do to make their firm healthier and to investigate what drives the recovery of a firm, this paper employed XGBoost models predicting the recovery of zombie firms.

These models find evidence that recovered firms are mainly characterized by their retained earnings, goodwill, tax expenses, rental expenses, and capital surpluses. Over the time period of this research, both goodwill and amortization have become increasingly important for recovering firms. On the other side of the spectrum, discontinuing operations, turnover, and several types of stock equity can hardly be considered to be very important in the recovery of a zombie firm.

Given these results, more in-depth qualitative research has been done to investigate the practical implications of the most important features. With this, it can be concluded that zombie firms who want to recover should increase their retained earnings by downsizing their workforce without decreasing wages or positive incentives. Whilst doing this, the focus should lie on the increase of the profit margins without price increases. Further, it can be concluded that zombie firms should decrease their rental expenses and fixed assets while focusing more on the intangible assets the firm possesses. This is important to increase capital surpluses and amortization. These measures also benefit from higher tax write-offs in the short term.

The recent increase in interest rates and the many different tax measures taken by governments can be interesting extensions to this research, as well as a more geographically focused dataset. This paper has taken a data scientific approach to use more innovative methods in finding the most important measures zombie firms can take to recover their firms. Most important features over time and their changing behavior in combination with a more qualitative discussion to investigate the practical measures that can be taken helped to gain an understanding of what characteristics recovering firms have and how this can be used to their advantage.

References

Acharya, V. V., Crosignani, M., Eisert, T., & Steffen, S. (2022). Zombie Lending: Theoretical, International and Historical Perspectives.

De Martiis, A., Heil, T., & Peter, F. J. (2021). Are you a Zombie? A Supervised Learning Method to Classify Unviable Firms and Identify the Determinants. A Supervised Learning Method to Classify Unviable Firms and Identify the Determinants (January 29, 2021).

Fukuda, S. I., & Nakamura, J. I. (2011). Why did 'zombie' firms recover in Japan?. The world economy, 34(7), 1124-1137.

Adalet McGowan, M., Andrews, D., & Millot, V. (2018). The walking dead? Zombie firms and productivity performance in OECD countries. Economic Policy, 33(96), 685-736.

Caballero, R., T. Hoshi and A.K. Kashyap (2008). 'Zombie lending and depressed restructuring in Japan', American Economic Review, 98, 1943–77.

Nakamura, J. I., & Fukuda, S. I. (2013). What happened to" zombie" firms in Japan?: Reexamination for the lost two decades. Global Journal of Economics, 2(02), 1350007.

Banerjee, Ryan, and Boris Hofmann (2018). "The rise of zombie firms: causes and consequences." In: BIS Quarterly Review September.

Choudhurry (2021). "Top XGBoost interview questions for Data Scientists'', Retrieved from: https://analyticsindiamag.com/top-xgboost-interview-questions-for-data-scientists/

Ho, Tin Kam (1995). 'Random Decision Forests' Proceedings of the 3rd International Conference on Document Analysis and Recognition, Montreal, QC, 14–16 August 1995. pp. 278–282.

Chabert, Jean-Luc (2012). A History of Algorithms: From the Pebble to the Microchip. Springer Science & Business Media. pp. 7–8. ISBN 9783642181924.

Nakamura, J. I. (2016). Japanese firms during the lost two decades: The recovery of zombie firms and entrenchment of reputable firms. Springer.

Banerjee, R., & Hofmann, B. (2020). Corporate zombies: Anatomy and life cycle. Economic Policy.

Yellen, J. L. (1984). Efficiency wage models of unemployment. The American economic review, 74(2), 200-205.

https://www.federalreserve.gov/

Chen, T, Guestrin, C. (2016) "XGBoost: A Scalable Tree Boosting System"

Grosse-Rueschkamp, B., Steffen, S., & Streitz, D. (2019). A capital structure channel of monetary policy. Journal of Financial Economics, 133(2), 357-378.

Appendices

Appendix A

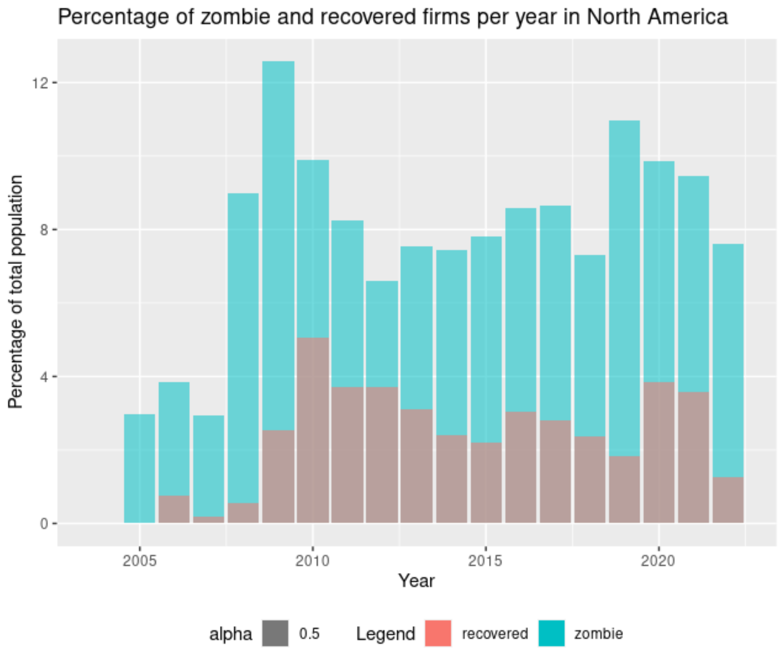

Graph A1: Percentage of zombie and recovered firms per year in North America, excluding Mexico. Note the bars are not stacked but two individual observations plotted together

Appendix B

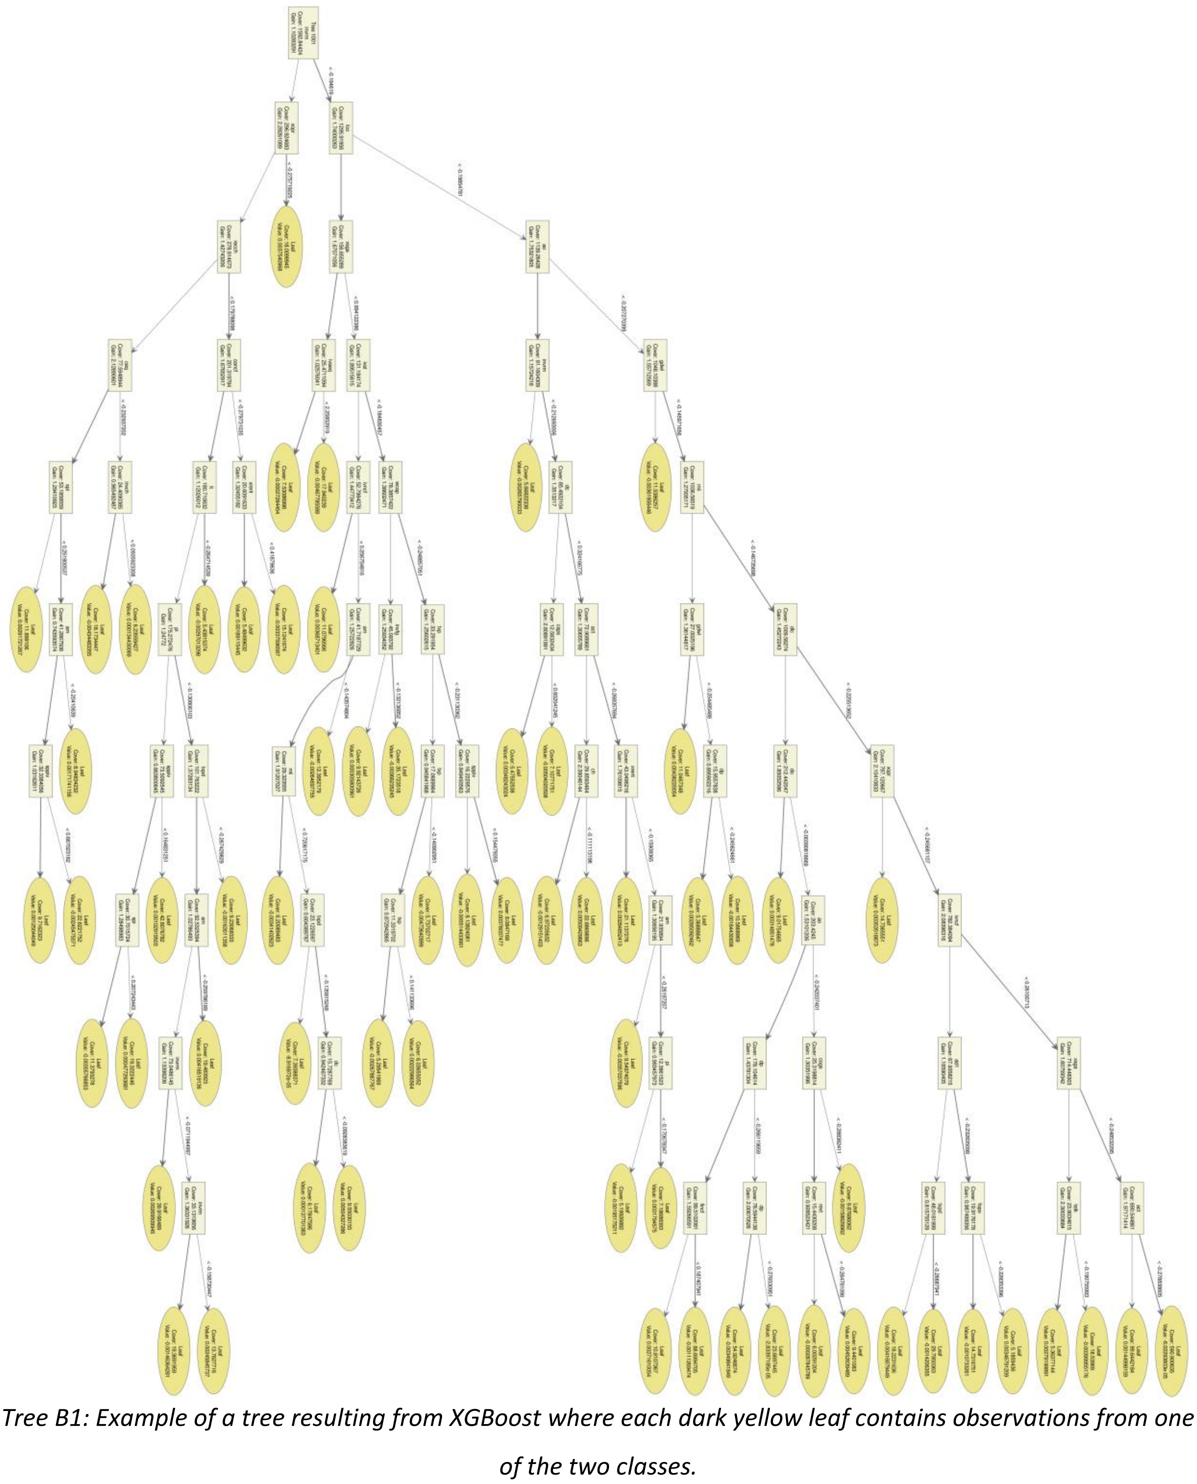

Tree B1: Example of a tree resulting from XGBoost where each dark yellow leaf contains observations from one of the two classes.

Appendix C

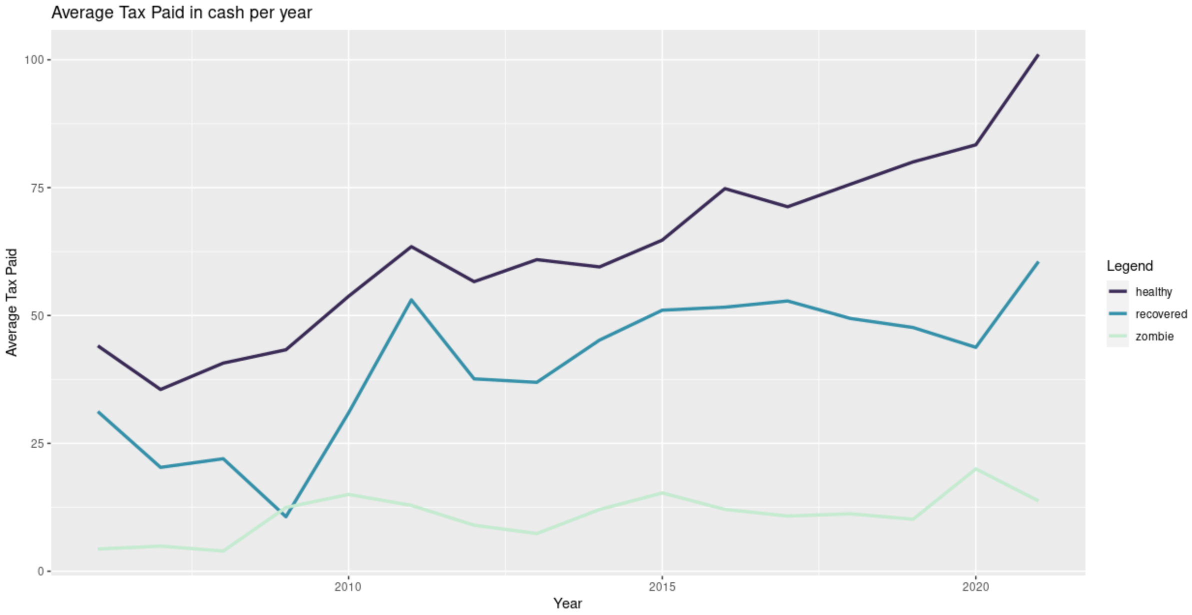

Graph C1: Average tax paid per year and per type of firm in millions

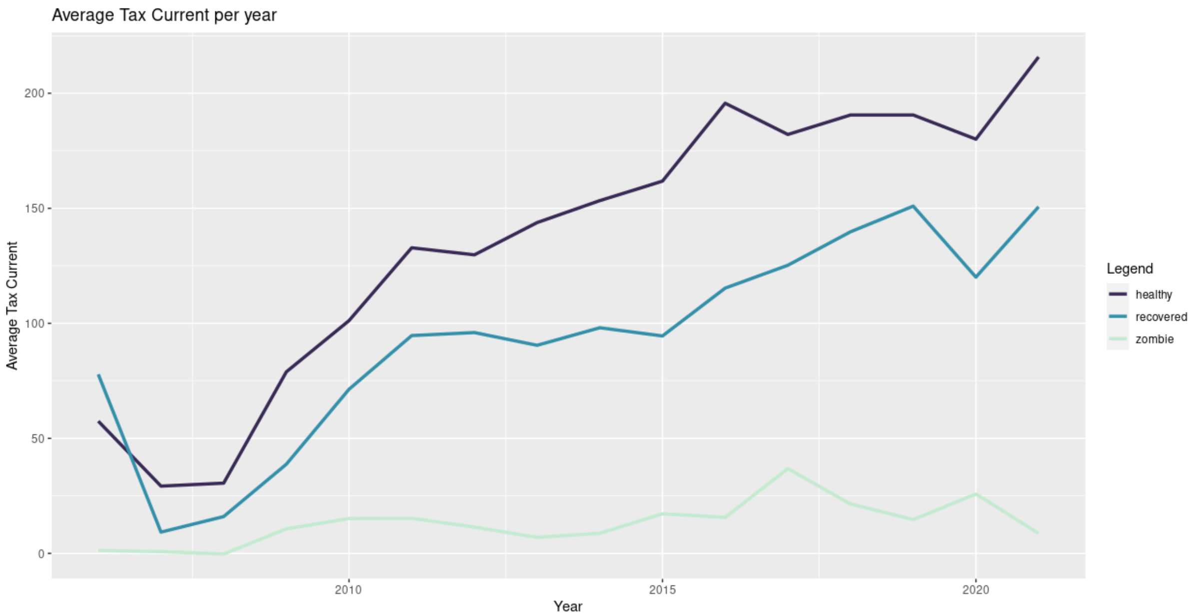

Graph C2: Average current tax per year and per type of firm in millions

Appendix D

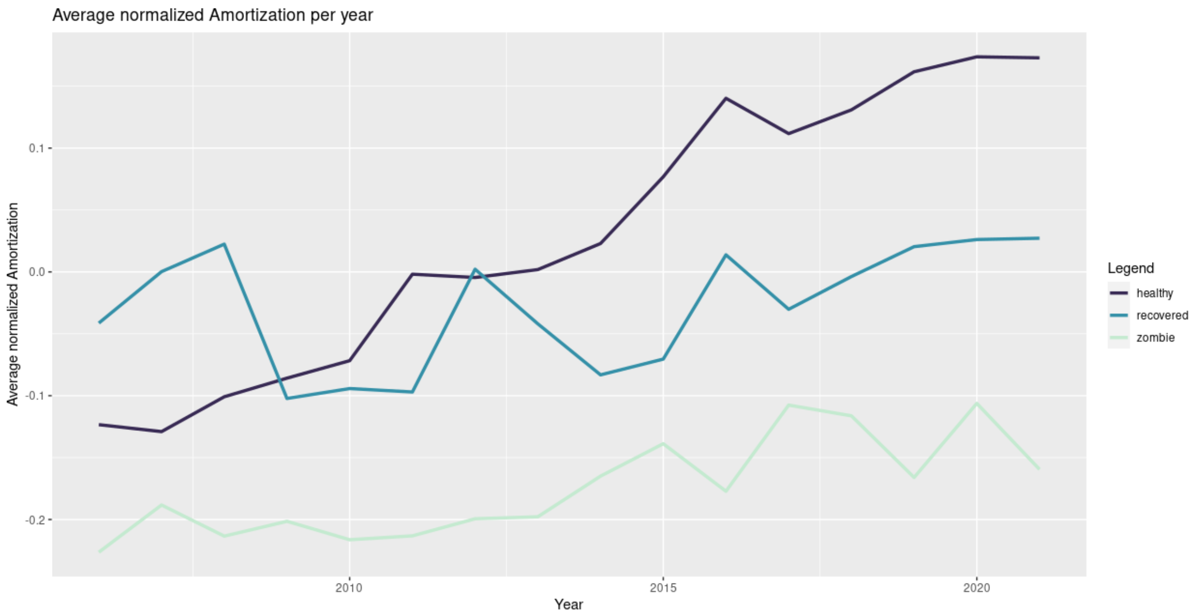

Graph D1: Average normalized amortization per year and per type of firm. Where 0.00 is the average amortization for all type of firms over the whole data period

Appendix E

2009		2010		2011		2012		2013	
dre	0.03996453	**txpd**	0.03737458	**re**	0.04039719	**re**	0.04067465	**re**	0.04200659
txpd	0.03796142	**re**	0.03564265	**txpd**	0.03349826	**txpd**	0.03278411	**txpd**	0.02742536
xrent	0.0225139	**caps**	0.02374101	**xrent**	0.02115869	**caps**	0.02255853	**xrent**	0.02236483
recch	0.02205801	**intpn**	0.02026424	**txc**	0.02075392	**ao**	0.02093945	**caps**	0.02180595
caps	0.02205458	**fopo**	0.020127	**caps**	0.01968159	**am**	0.02069943	**am**	0.01993793
ao	0.02175577	**xrent**	0.01924196	**dp**	0.01917671	**intpn**	0.01946871	**ch**	0.01992339
idit	0.02125566	**txc**	0.01920503	**idit**	0.01895393	**xrent**	0.01911821	**txc**	0.01985356
dp	0.02074946	**idit**	0.01888357	**fopo**	0.01882371	**ebit**	0.0185364	**intpn**	0.01950575
ivncf	0.01995192	**xacc**	0.01870565	**dd1**	0.01855922	**recch**	0.01818902	**recch**	0.01934492
spi	0.01909349	**ao**	0.01834375	**recch**	0.01815344	**invch**	0.01807707	**intan**	0.01917758
fopo	0.01901908	**recch**	0.01800744	**intpn**	0.01814173	**dd1**	0.01800281	**idit**	0.01876594
txdi	0.01879106	**xsga**	0.0177548	**am**	0.01806267	**lo**	0.01724505	**gdwl**	0.01847723
xacc	0.01866103	**gdwl**	0.0176575	**aoloch**	0.01798927	**xacc**	0.01703877	**ebit**	0.018426
am	0.01857007	**am**	0.01720903	**xsga**	0.01742207	**txc**	0.01667462	**ao**	0.01837607
xsga	0.01787504	**dp**	0.01679193	**ao**	0.0173762	**aoloch**	0.01666117	**fiao**	0.01734688
aoloch	0.01727861	**aoloch**	0.01655092	**gdwl**	0.01670421	**xsga**	0.01620019	**dd1**	0.01663821
sppiv	0.01713981	**sppiv**	0.01629989	**invch**	0.01653873	**intan**	0.01617872	**spi**	0.01651454
ap	0.0167263	**ch**	0.01628471	**ebit**	0.01642201	**fopo**	0.01587725	**invch**	0.01628144
invrm	0.01642593	**invwip**	0.01588677	**xacc**	0.01633427	**pi**	0.01569221	**aoloch**	0.01620234
rectr	0.01632077	**ebit**	0.01562827	**intan**	0.01591332	**ap**	0.01557176	**dp**	0.01584275

2014		2015		2016		2017		2018		2019	
re	0.03958858	re	0.04009628	re	0.03999637	re	0.03120144	re	0.03941319	re	0.03745791
txpd	0.0363478	txpd	0.03427541	txpd	0.02779392	gdwl	0.02400393	txpd	0.02673849	gdwl	0.02877887
txc	0.02210469	gdwl	0.02386434	ao	0.0261262	caps	0.0219479	gdwl	0.02281791	caps	0.0245784
xrent	0.02194705	xrent	0.02219638	gdwl	0.02315027	txpd	0.02150961	caps	0.02267423	intpn	0.02356199
intan	0.02175263	intpn	0.0221642	am	0.0229783	ao	0.02144332	intpn	0.02155735	am	0.02291403
gdwl	0.02103934	ao	0.02176558	intpn	0.0228975	idit	0.02100466	xacc	0.02124729	txpd	0.02209115
intpn	0.02100002	caps	0.02117314	txc	0.02075399	txc	0.02040133	xsga	0.02037446	xsga	0.02203486
caps	0.02059682	idit	0.02073098	caps	0.02072869	intpn	0.01989916	ao	0.01987497	xrent	0.02122109
am	0.02002201	am	0.02018312	idit	0.01915931	intan	0.01958151	am	0.01986943	idit	0.02081822
ao	0.01865098	xsga	0.01924864	intan	0.0187009	recch	0.01928357	spi	0.01976638	recch	0.02033847
xsga	0.01805947	intan	0.01895345	recch	0.01856574	xsga	0.01915092	idit	0.01966332	ao	0.02006012
recch	0.01797977	txc	0.01891632	xrent	0.01843975	am	0.01881914	xrent	0.01952835	intan	0.01976173
pi	0.01683855	fiao	0.01833613	ebit	0.01784253	aoloch	0.01805836	recch	0.01889069	spi	0.01958944
fiao	0.01636365	ch	0.01826028	capx	0.01712927	spi	0.01670389	intan	0.01826934	xacc	0.01951771
aoloch	0.01628372	aoloch	0.01825591	xsga	0.01680335	ebit	0.01608219	aoloch	0.01737542	fiao	0.01925255
ppent	0.01618391	recch	0.01811107	spi	0.01674405	xacc	0.01594942	dd1	0.01681918	aoloch	0.01896364
ebit	0.01608329	ebit	0.01677299	aoloch	0.01665281	fiao	0.01589284	invwip	0.01669433	invfg	0.01738843
invwip	0.01596635	spi	0.01653985	invrm	0.01647401	xrent	0.01584097	fopo	0.01635426	invch	0.01738533
ap	0.01573267	invwip	0.016449	dd1	0.01615399	invwip	0.01548596	txc	0.01621552	dd1	0.01722052
spi	0.01569943	invrm	0.01619802	xacc	0.01604761	invfg	0.01541527	lo	0.01583797	ivaco	0.01368742

Table E1: The top 20 most prominent features for XGBoost per starting year. Importance is measured as gain. Abbreviations can be found in main text or at Compustat's website.

Appendix F

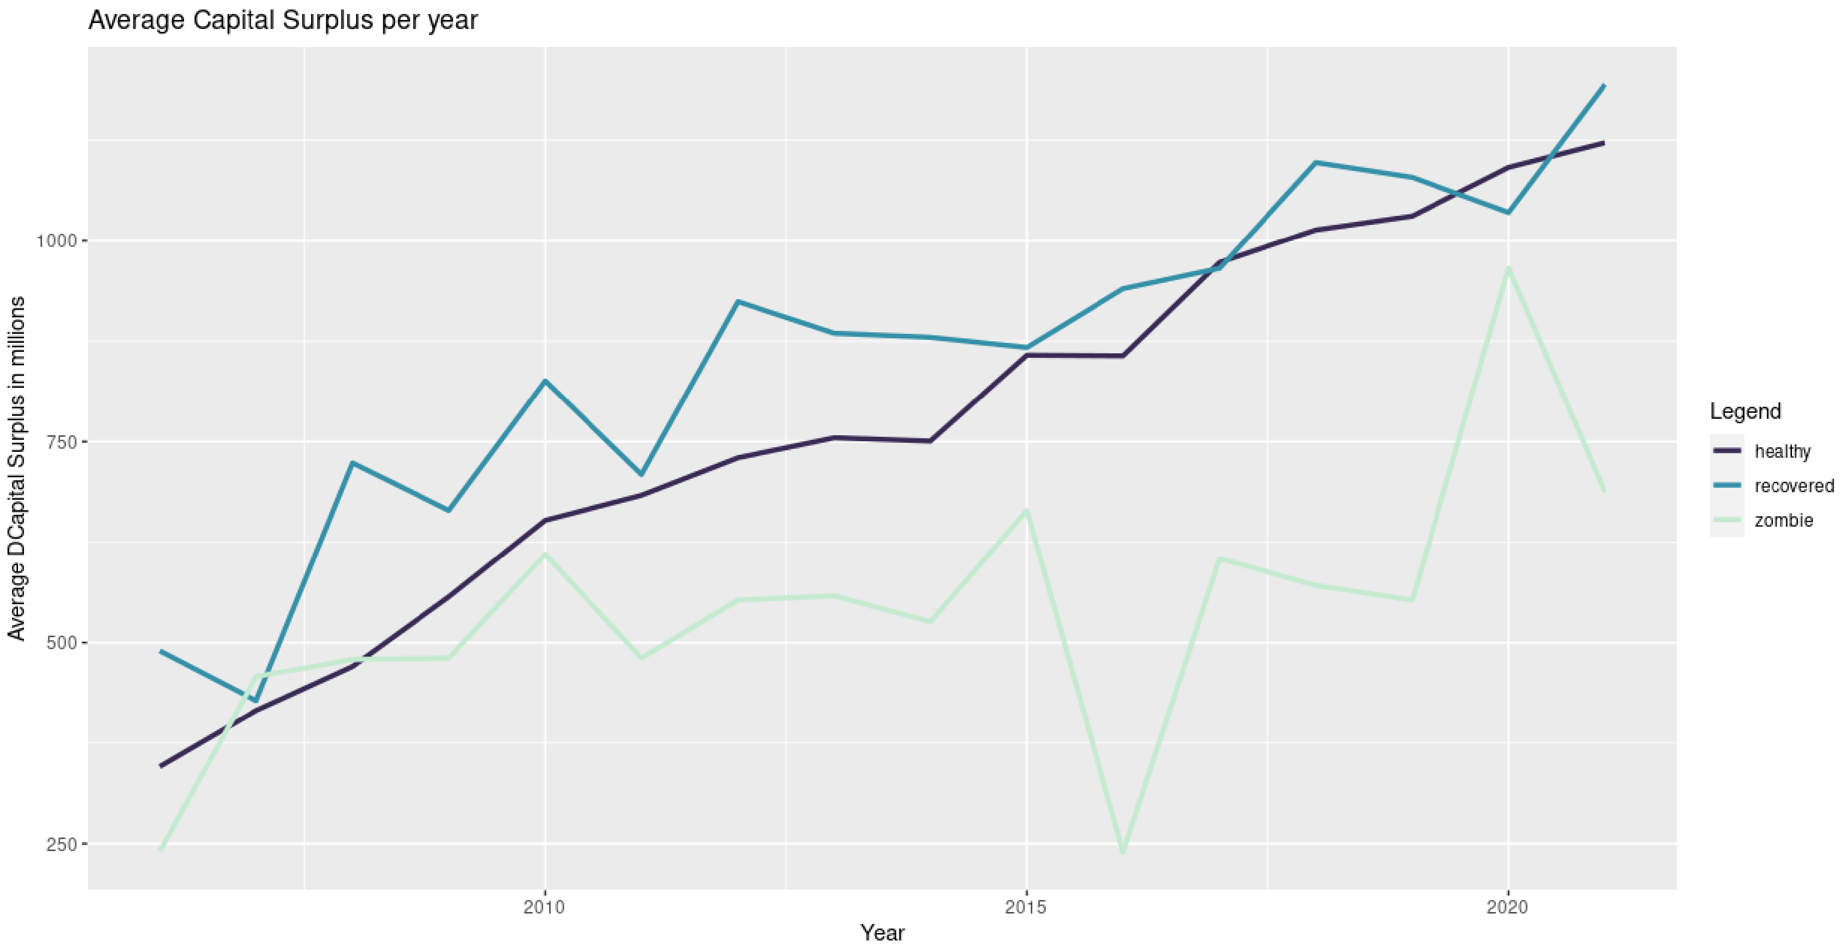

Graph F1: Average Capital Surplus in millions per year and per type of firm.