Predicting Takeover Targets*

- A Machine Learning Approach

Grim Sjöberg† and Dag Tuvberger‡

Stockholm School of Economics

May 16, 2020

Abstract

The possibility of predicting takeover announcements for investment purposes has been widely researched, with varying results and conclusions. This paper builds on previously conducted studies and applies a machine learning approach to the research problem, through the random forest classifier by Breiman (2001). The classifier is known to perform well on imbalanced datasets, which is a key characteristic of takeover data. Furthermore, this paper applies a robust research methodology by extending the time period studied and replacing the choice of an optimal classification cutoff probability with a range of likelihoods. With this range of probabilities, it investigates the trade-off between predictive accuracy and size of portfolios. We conclude that the addition of the random forest classifier does not bring statistically significant superior results compared to previous models. However, our findings show that it may indeed be possible to earn statistically significant abnormal returns through an investment strategy of takeover target prediction. As these findings are not consistent over time, further research is warranted.

Keywords: Takeover, Prediction, Random Forest, Machine Learning

JEL Classification: G34 (Mergers; Acquisitions; Restructuring; Corporate Governance), G11 (Portfolio Choice; Investment Decisions), C53 (Forecasting and Prediction Methods; Simulation Methods)

*We would like to thank Michael Halling, Associate Professor at the Department of Finance at the Stockholm School of Economics, for his insightful comments and support throughout our work. Furthermore, we extend a thank you to Ronak Shah for his valuable advice concerning the construction of our prediction models. This paper was written as a master thesis at the Department of Finance at the Stockholm School of Economics.

†E-mail: 23307@student.hhs.se

‡E-mail: 23395@student.hhs.se

Contents

List of Tables

List of Figures

1 Introduction

"*It is difficult, if not impossible, for the market to identify future target firms*"

- Ruback and Jensen (1983)

Several papers have tried to use statistical methods and models to predict takeover targets, with notably different results and conclusions. The motivation to do so is mainly based on the possibility of earning abnormal returns by trading on this information. Ruback and Jensen (1983) and more recently Andrade et al. (2001) show that shareholders of target firms experience large and statistically significant abnormal returns during multiple event windows centered around the announcement date of acquisition.[1] The possibility of being able to predict these takeovers, and more specifically the time of announcement, would therefore compose a compelling and lucrative strategy for any investment manager.

We build on previously conducted studies by using more recent statistical techniques to predict takeover announcements. Previous papers have mainly relied on logistic regression to predict these announcements, a standard statistical technique when dealing with binary outcomes. We apply a machine learning algorithm, more specifically the random forest algorithm, to research whether this application could improve prediction results.

Furthermore, we seek to enhance robustness of our results through a series of methodological changes compared to previous studies. Firstly, we extend the time window studied from one individual year to three consecutive years. Secondly, we use a distribution-based range of cutoff probabilities for when to classify a firm as a future takeover target, rather than attempting to find an optimal value. This allows us to study the trade-off between predictive accuracy and size of the portfolio over different cutoffs and the impact from this trade-off on returns. In addition to these two specific methodological changes, we seek to throughout the paper adhere to a method that relies only on information available to the decision maker ex ante an investment decision. In this way, our ambition is to produce results with practical implications.

We find that it is possible to construct prediction models with predictive value in the area of takeover target prediction. However, the predictive value is small. The addition of the random forest algorithm does not significantly improve predictive value. We use the predictions to build portfolios with a holding period of one year, where firms experiencing takeover bids are sold on the day of bid announcement, to evaluate whether our predictions

[1] For more evidence, see e.g. Langetieg (1978), Jarrell and Poulsen (1989), Schwert (1996) and Goergen and Renneboog (2004)

could be the basis for a successful investment strategy. In this evaluation, we see significant abnormal returns in a minority of the portfolios formed. We conclude that the viability of this strategy can not be rejected, but warrants further research.

The disposition of this paper is as follows. We begin with describing the research background surrounding takeover announcement returns and takeover target prediction. In section two, we then present the data sample used for this study and outline the methodology of our study extensively. In section three, we present the empirical results of our study and finally in section four results are discussed and conclusions are drawn.

1.1 Background

The characteristics of abnormal returns for both target and acquiring shareholders surrounding a takeover announcement is a well studied subject. Focusing solely on target shareholders, there are numerous studies showing that takeovers are beneficial, consistently describing large premiums being paid by the acquirer. Mandelker (1974), Ruback and Jensen (1983) and Bradley et al. (1983) all write about abnormal returns attributable to target shareholders in the event of a takeover, with gains ranging between as much as 20-30% for successful mergers and tender offers. According to data of all US deals from 1980 to 2005, the average price reaction on the day of announcement was 15% for target firms (Betton et al. (2008)). Schwert (1996) looks at an event window for successful deals only starting 42 days before the announcement ending 126 after, concluding that target shareholders earn on average cumulative abnormal returns of 30.1% during this period. Bradley et al. (1983) also distinguish between unsuccessful and successful takeover announcements, concluding that even though also unsuccessful bid announcements experience positive abnormal returns on the day of announcement, the long term performance of target firms vary significantly. The study finds a negative cumulative abnormal return over a one year period for firms that have not been taken over following an unsuccessful takeover announcement.

Concerning the prediction of takeover targets, Palepu (1986) is one of the more influential papers. The paper pinpoints a number of methodological shortcomings carried out by earlier studies claiming impressive abilities to predict targets 6 to 12 months before the announcement of takeover. Palepu (1986) corrects for these shortcomings and conducts a new empirical study, showing that takeover prediction indeed is difficult to master, and that "predictions reported by the earlier studies are overstated". He further addresses the discrepancy between the predictive accuracy of previously conducted studies and the stock market's ability to predict future takeover targets. Seeing limited movement in stock prices even days before the announcement clearly shows the market is unaware of the takeover before its announcement. In other words, if predictions were as accurate as stated in studies criticized by Palepu (1986), wouldn't investment managers reap these profits until there weren't any left?

Using a logistic regression model with nine independent variables, Palepu (1986) finds his model statistically significant, but with small explanatory power. The problem is that even though the model successfully predicts a high proportion of actual targets, it also misclassifies a significant number of non-targets as targets. These misclassifications dilute the portfolio return, making it impossible to earn significant abnormal returns by trading on the modelled predictions. His investment strategy is to define an optimal cutoff probability and then taking a long position in companies whose predicted takeover likelihood exceeds this cut-off. By looking at the modelled distribution probabilities for target and non-target firms in the training sample, the optimal cutoff probability is found where the probability distribution of target and non-target firms intersect. Above this specific probability, the probability function for target firms is higher than for that of non-target firms.

With an optimal cut-off probability of 11.2%, predictions in the test sample of Palepu (1986) amounts to 625 targets and 492 non-targets. However, actual targets and non-targets were in fact 30 and 1,087, respectively. Using an equally weighted portfolio over a period of 250 trading days the CAR is measured at -1.6%, more negative than the CAR for predicted non-targets at -1,5%. However, the correctly predicted 24 targets experienced a CAR of 21% during the same period, while the 6 targets classified as non-targets showed CAR of 36.2%, potentially indicating that more unexpected takeover announcements generate better returns to target shareholders.

Powell (2001) extends the work of Palepu (1986) by defining a portfolio selection criterion based on maximizing the percentage of actual targets in the portfolio, but otherwise rely on similar methodology. Instead of minimizing the total number of misclassifications like Palepu (1986), he creates ten portfolios after sorting firms in order of takeover likelihood and then base his decision on the portfolio with the highest proportion of targets. Powell (2001) also comes to the conclusion that forming an investment decision based on takeover prediction models does not yield significant abnormal returns.

Furthermore, Brar et al. (2009) build on Palepu (1986) by adding momentum, trading volume and market sentiment as independent variables. They argue these variables do a better job in capturing the timing of the acquisition announcement. Brar et al. (2009) find, contrary to previous studies, that their prediction model can indeed be used to build portfolios that earn significant abnormal returns.

2 Data & Methodology

In this section we describe the data sample used in our study. We further describe the prediction models employed in detail and how we apply the predictions to make investment decisions. The study is based on recording a number of independent variables for firms in a twelve-month estimation window and consequently observing whether the firms receive a takeover bid in a twelve-month observation window. This means all observations are paired twelve-month observations with data from both an estimation window and a subsequent observation window.

2.1 Data Sample

To train and evaluate the prediction models, we use data on all index constituents of the S&P Super Composite index between the years 2010-2019.[2] Firms within the financial services industry have been dropped to ensure the peculiar characteristics of banks' financial ratios do not impact our models. The S&P Super Composite is a broad market value-weighted index with, at any point in time, 1500 constituents. The index includes all stocks in the S&P 500, S&P 400 and S&P 600, representing around 90% of the total market capitalization of all US equities. Given the relative rarity of takeover bids, we believe the large number of companies help satisfy the need for data to successfully train our prediction models.

We start by splitting our data into six subsets, three used for training and three for testing. The training and testing datasets are then paired, so that one training dataset is used to train a model which is then used for predictions on its paired test dataset. Figure 1 shows an illustration of how the different datasets are split and paired.

Our rationale for splitting the data this way is to mimic a real-world situation where one would have access to earlier data for training and then implement the trained classifiers on observations later in time, while also retraining the classifier each year to incorporate new information.

2.1.1 Potential Independent Variables

To find appropriate independent variables to use in our prediction models, previous research into the area of takeover target prediction is consulted. Earlier papers such as Palepu (1986) and Brar et al. (2009) list the following hypotheses on why companies receive takeover bids[3]:

[2]Our dataset is gathered using databases from the Wharton Research Data Services and the M&A specific database SDC Platinum.

[3]see table 1 for a summary of these hypotheses, their independent variables, and corresponding impact on the likelihood of acquisition

Figure 1: Subsets of data

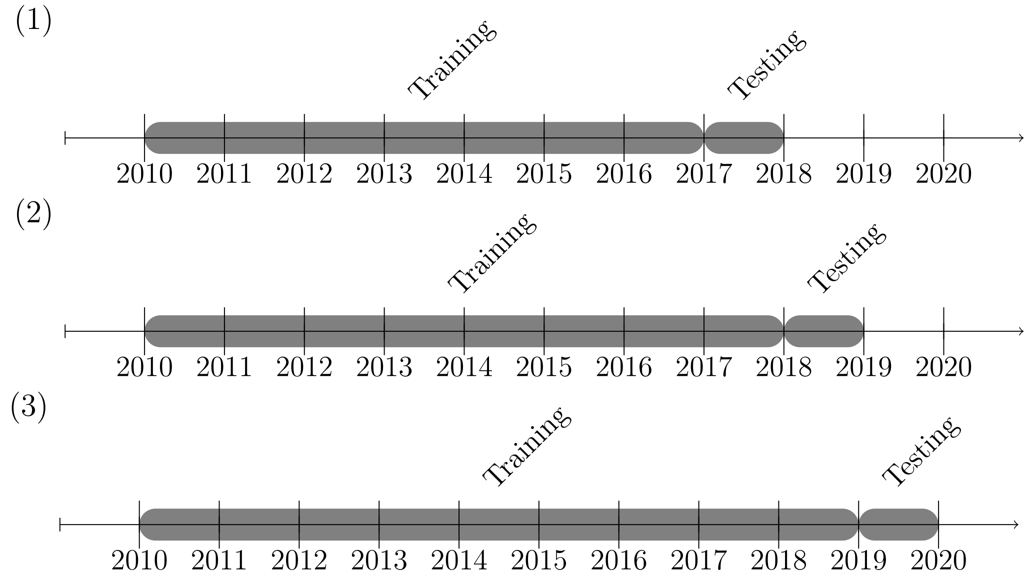

This figure shows how the data sample is divided into different test and training datasets

1. *Inefficient management hypothesis: Under-performing firms are likely targets of acquisition.*

Acquisitions take place to replace management who fail at maximizing the market value of a firm. Relatively under-performing firms are therefore more likely targets of acquisition. Five variables are suggested to test this hypothesis:

- *Average excess return*
- *Net margin*
- *Operating margin*
- *Sales growth*
- *Return on equity*

2. *Growth-resource mismatch hypothesis: firms with significant growth opportunities but low financial resources or vice versa are likely targets of acquisition.*

A growth-resource mismatch means there is a mismatch between a firm's growth opportunities and its financial resources. Resource-rich firms experiencing low growth are underperforming. However, also high-growth firms with low resources are attractive targets as they provide the acquirer with ample growth opportunities at a relatively low price. The following dummy is proposed to test for this hypothesis:

- *Growth-resources dummy*

3. *Industry disturbance hypothesis: Firms in industries with high M&A activity are likely targets of acquisition.*

Acquisitions are clustered by industry and time. Thus, firms in industries which have recently seen acquisitions are more likely to be targets. A dummy variable is suggested to test for this hypothesis:

- *Industry disturbance dummy*

4. *Firm size hypothesis: Smaller firms are likely targets of acquisition.*

Acquisitions entail transaction costs and these costs increase with the size of the target (e.g. integration costs post merger and takeover defense cost). Larger targets will therefore have fewer bidders who can sustain these costs and thus smaller firms should be more likely targets of acquisition. We test this hypothesis using the following variables:

- *Net sales*
- *Market capitalization*
- *Total assets*

5. *Market-to-book hypothesis: firms with low market value of equity relative to book value of equity are likely targets of acquisition.*

Firms with a low market-to-book ratio are thought of as cheap or undervalued and are therefore more likely to be targeted.

- *Market-to-book ratio*

6. *Price-earnings hypothesis: Firms with a low PE-multiple are likely targets of acquisition.*

Bidders with a high price-earnings ratio seek to acquire targets with a low price-earnings ratio because of the belief that the stock market values the combined entity at the bidder's higher price-earnings ratio. Thus, firms with lower price-earnings ratio are more likely to be targets of acquisition.

- *Price-earnings ratio*

Brar et al. (2009) further extend the hypotheses above by suggesting that firms with higher leverage and higher short-term momentum in their stock prices are more likely targets. Also measures for liquidity are included in their study with the expectation that relatively less liquidity increases takeover probability.

In table 1, we summarize independent variables suggested in both papers. We also present the expected sign of regression coefficients for these variables, i.e. how they correlate with companies' acquisition likelihood in a given period. The computation of independent variables follows previous literature (Palepu (1986) and Brar et al. (2009)) and is described in table 12.

Table 1: Independent variables

Hypothesis	Variable(s)	Expected sign
Inefficient management	Average excess return (AER)	-
	Net margin	-
	Operating margin	-
	Sales growth	-
	Return on equity (ROE)	-
Growth-resource mismatch	Growth-resources dummy (GRDUMMY)	+
Industry disturbance	Industry disturbance dummy (IDUMMY)	+
Firm size	Net sales	-
	Market capitalization	-
	Total assets	-
Undervaluation	Market-to-book value (MTB)	-
	Price-earnings ratio	-
	Earnings yield	-
	Dividend yield	-
Momentum	12-month price momentum	+
Leverage	Short-term leverage	+
	Long-term leverage	+
Liquidity	Liquidity	-
	Cash-to-capital	-

This table present hypotheses discussed by Palepu (1986) and Brar et al. (2009) on why firms are being taken over, variables that correspond to each hypothesis and their expected regression coefficient sign on the likelihood of takeover. Computation of variables is described in table 12.

2.2 Methodology for Prediction

In our study, we utilize two different prediction models; a logistic regression model and a random forest classifier. We start by using the more commonly used model of logistic regression, mimicking its application in previous papers, to create a baseline for predictive power. The logistic regression itself is oftentimes referred to as the standard approach for binary classification, especially in scientific fields where focus lies more on explanation rather than prediction (Couronné et al. (2018)). Secondly, we apply a random forest classifier to the prediction problem. This model puts more emphasis on prediction rather than explanation, a trait beneficial for a potential investment manager looking to trade on the predicted outcomes.

The disposition of this this section is as follows. We start by describing the theoretical foundation of the logistic regression model in detail and furthermore our specific application of it. Next, we introduce a theoretical framework describing the advantages of using a random forest classifier in the problem of predicting future takeover targets. We describe the random forest algorithm in detail and finally present a methodology for comparing the two models.

2.2.1 Binomial Logistic Regression Model

Binomial logistic regression is a standard approach to binary classification tasks. In this specific regression analysis technique, the dependent variable has a binomial distribution. Therefore, the methodology is sometimes referred to as a binary regression. The technique fits a linear model between the independent variables and the log-odds of the probability of a positive outcome in the dependent variable. Log odds is defined in equation 1. The logarithm of the odds is used as the probability only ranges between 0 and 1.

$$Log\ odds = log \frac{p}{1-p} \tag{1}$$

Denoting the conditional probability of a certain binary outcome $P(Y = 1|X_1, ..., X_n)$ as $p(x)$, the logistic regression model can be written as:

$$log \frac{p(x)}{1-p(x)} = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + ... + \beta_n x_n \tag{2}$$

$$\frac{p(x)}{1-p(x)} = e^{\beta_0 + \beta_1 x_1 + \beta_2 x_2 + ... + \beta_n x_n} \tag{3}$$

$$p(x) = \frac{e^{\beta_0+\beta_1 x_1+\beta_2 x_2+...+\beta_n x_n}}{e^{\beta_0+\beta_1 x_1+\beta_2 x_2+...+\beta_n x_n}+1} = \frac{1}{1+e^{-\beta_0+\beta_1 x_1+\beta_2 x_2+...+\beta_n x_n}} \tag{4}$$

Lastly, the binomial logit model (equation 5) describes the relationship between firms' characteristics and their likelihood to be acquired in the given period.

$$p(i,t) = \frac{1}{1+e^{-\beta x(i,t)}} \tag{5}$$

where $p(i,t)$ is the probability that firm i will be acquired in period t, $x(i,t)$ a vector of our independent variables, and β a vector of coefficients estimated by the model.

For our application of the logistic regression model, we choose to exactly replicate Palepu (1986) in the choice of variables. By doing so, we create a baseline model to compare the random forest model to. This includes fitting four regression models and finally applying model two of these for predictions. Model two includes the following independent variables:

- *Average excess return*
- *Growth-resource dummy*
- *Growth*
- *Liquidity*
- *Leverage*
- *Industry disturbance dummy*
- *Size (Market capitalziation)*
- *Market-to-book ratio*
- *Price-earnings ratio*

We train the logistic regression model on a random sample of two thirds of the three training data samples. The reason behind this is that the second prediction model used, the random forest classifier, puts aside a third of training data to construct an out-of-bag dataset. To allow for the comparison between logistic regression and random forest, we must ensure the two classifiers are trained on the same amount of data. Therefore, we randomly sample two thirds of the training datasets described and use for training on the logistic regression, while keeping one third as an out-of-bag dataset. This further allows us to report model performance metrics also for the years 2010-2016, when we would otherwise have no test data.

2.2.2 Random Forest Model

The problem of predicting future takeover targets is a highly imbalanced classification problem as it is far more uncommon for a company to receive a takeover bid in a given time period than for it not to. The random forest algorithm has in previous benchmarking studies been reported to perform better than logistic regression on imbalanced data. Brown and Mues (2012) compares a number of classifiers including random forests and logistic regression on five credit scoring datasets which all exhibit class imbalance. The random forest classifier consistently performs better than the logistic regression model and the difference in performance increases with the degree of class imbalance. Another benchmark study between random forest and logistic regression on imbalanced data is Muchlinski et al. (2016), which showed that Random Forests is significantly better than logistic regression at predicting the rare event of civil war onsets.

Outside of class imbalance problems, random forests have also been described as generally having stronger predictive power than logistic regression. Couronné et al. (2018) performs a large-scale benchmarking study on not less than 243 datasets (from various scientific disciplines) and shows that random forests outperform logistic regression in 69% of the datasets.

In conclusion, these recent findings make us ponder if the application of random forests to the takeover prediction problem could yield favorable results, exploitable in an investment management situation.

We use the random forest model as developed by Breiman (2001). The model utilizes decision trees, a simple machine learning algorithm that uses a tree-like structure to perform classification or regression tasks. In each leaf (called node) of the tree, independent variables are evaluated to train a decision rule that most accurately separates the data with respect to the dependent variable. Each node is then followed by two branches in the decision tree, in turn followed by new nodes where new independent variables are evaluated and thus a tree-like structure is created (figure 2).

The random forest model is an ensemble of many decision trees where, for classification tasks, the modal class of the ensemble of trees is the final predicted class of the random forest. For regression tasks, the average value of the ensemble is instead used. In this way, random forests operate on classification tasks similar to a committee voting for a particular outcome. Random forests reduces generalization error (i.e. they overfit less) compared to single decision trees. This is due to two distinct features of the algorithm. Firstly, each tree in the random forest is trained on different samples of the data using a sampling-with-replacement technique (bootstrapping). Secondly, each node in the ensemble trees are fed a predetermined number of randomly chosen independent variables to form its decision rule on (random subspace method) rather than being trained on all independent variables at once.

Figure 2: Decision tree sample structure

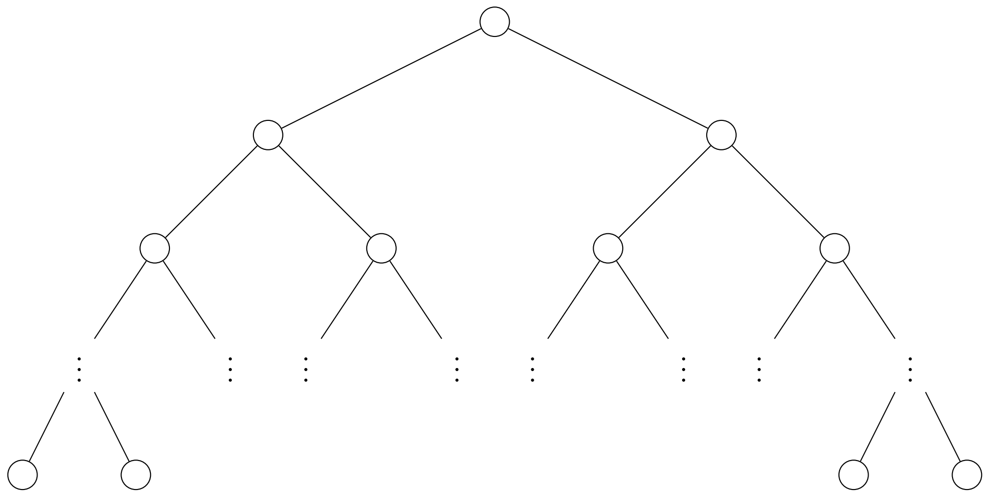

This figure illustrates the mechanism of a decision tree algorithm. Each round circle represents a node, where the algorithm evaluates predictive variables and sets a split value for each variable. The split value is set with the objection of splitting the dataset into two groups as homogeneous as possible with respect to the dependent variable.

Both these distinct features serve to reduce correlation among the trees in the ensemble and, consequently, to reduce the generalization error of the model. The ensemble trees are, according to Breiman (2001), grown to maximum size and are not pruned (removal of nodes).

In Breiman's random forest model for classification tasks, the CART methodology (Breiman et al. (1984)) is used to split data in each tree node. Data is split at each value an independent variable is observed at in the sample. A performance metric for each split is then calculated and the best performing split value is chosen for that particular independent variable. The process is then repeated for all independent variables randomly chosen (according to the random subspace method) in that node. The performance metric used in the CART methodology is Gini Impurity:

$$Gini\ Impurity = 1 - \sum_{i=1}^{C} (\rho_i)^2 \tag{6}$$

where C is the number of classes and ρ is the probability of misclassification of a class. Gini Impurity is calculated in the sub-nodes for each split in a parent node. Gini Impurity in the sub-nodes are then weighted to a combined figure for each split. Subtracting the sub-node weighted Gini Impurity from the parent node Gini Impurity, the decrease in Gini Impurity can be obtained. The split value with the highest decrease in Gini Impurity is then selected.

In the previously described bootstrapping process of the random forest algorithm, about two thirds of the training data set are used. The final third (called out-of-bag or OOB data) is saved and used to calculate model predictive performance measures. We use the out-of-bag data in tuning the parameters of the model. The main parameters of the Breiman (2001) random forest model are the number of decision trees in the ensemble and the number of independent variables selected at random in each tree node. Using the out-of-bag data and optimizing for the performance measure ROC AUC (see section 2.2.3), we apply an iterative approach to tuning the values of these parameters by continuously and in small steps increasing the values and studying the change in ROC AUC values. This approach is repeated for each of the three training datasets. We decide to produce 600 decision trees for the first random forest model (trained on the dataset 2011-2016), 300 for the second and 1000 for the third (see figure 8). For the number of independent variables selected at random in each node, a value of four is selected for the first training dataset, four for the second and two for the third dataset (see figure 9).

The out-of-bag portion of the training data can also be used to calculate the importance of different independent variables in the prediction model. This is done by, for each ensemble tree, calculating the sum of decrease in Gini Impurity for each independent variable. Then, a mean for each independent variable over all ensemble trees is calculated. This mean can then be compared to study the relative importance of independent variables.

In our application of the random forest model, we include all variables presented in section 2.1.1 without discrimination. Using the random forest's mechanism of calculating variable importance described previously, we can then draw conclusions on which variables carry the most predictive power.

2.2.3 Evaluating Models

To evaluate the logistic regression and random forest models, we use the area under the curve of the Receiver Operating Characteristic (ROC AUC). The reasons for this are two-fold. Firstly, the problem of predicting future takeover targets deals with highly imbalanced data, as it is more uncommon for a firm to receive a takeover bid than for it not to. Secondly, we seek an evaluation metric that is independent of cutoff probabilities, so that we analyze the models' predictive powers over a range of cutoff probabilities instead of being constrained to a single one. The ROC AUC works well on data with class imbalance and is independent of probability cutoffs (Bradley (1997)). To understand ROC and ROC AUC it is helpful to first study a binary classification confusion matrix (table 2).

Table 2: Classification confusion matrix

	Actual target	Actual non-target
Predicted target	True Positive	False Positive
Predicted non-target	False Negative	True Negative

This table illustrates a generic confusion matrix for a binary classification problem. More specifically, it presents the four categories an observation can be separated into after a classification model (with a set cutoff probability) has been applied to it. From this matrix, a number of classification performance measures can be calculated.

ROC is defined as a plot of the true positive rate (7) of a classification model against the false positive rate (8) at different probability cutoffs.

$$True\ Positive\ Rate = \frac{True\ Positives}{True\ Positives + False\ Negatives} \tag{7}$$

$$False\ Positive\ Rate = \frac{False\ Positives}{False\ Positives + True\ Negatives} \tag{8}$$

The ROC AUC is subsequently the integral of the curve produced by this methodology. The metric can be interpreted as the probability of a classification model to assign a higher estimated probability of positive class to a true positive observation than to a true negative observation. Thus, a ROC AUC of 0.5 implies a model which performs on par with random guessing (50% chance to assign a higher probability of positive class to a true positive than to a true negative). Higher (or lower) ROC AUC values imply a model with predictive power.

A ROC AUC of 1 signifies a perfect model.

To test whether the ROC AUCs of the logistic regression model and random forest model differ at a statistically significant level, we employ the DeLong test (DeLong et al. (1988)). The DeLong test is a non-parametric method for comparing ROC AUCs for paired ROCs. Paired ROCs are defined as ROCs calculated for two variables observed in the same sample. As we train our two classifiers on the same dataset (albeit with a difference in variables employed), the DeLong test is appropriate to use.

Apart from the ROC AUC, we apply one more metric in evaluating our prediction models. This is the positive predictive value (PPV):

$$Positive\ Predictive\ Value = \frac{True\ Positives}{True\ Positives + False\ Positives} \tag{9}$$

Applied to the problem of predicting takeover targets, the positive predictive value is equal to the ratio of actual targets in the predicted portfolio. We aim to maximize the proportion of actual targets in our trading portfolios, trying to make sure abnormal returns are not diluted by firms not experiencing a takeover bid. Thus, the higher the positive predictive value of our prediction models, the better. We therefore analyze this metric ex post to evaluate the models' capabilities to construct portfolios with a high concentration of actual targets (true positives).

2.3 Methodology for Investment Strategies

In this section, we start with a description of the methodology used for evaluating risk adjusted performance of portfolios in this study. We then give a brief overview of the reasoning behind our strategy of investing in firms subject to takeover announcements. Lastly, we combine the investment strategy with our prediction models, and show how we intend to use the estimated takeover likelihood from these models to build portfolios of firms believed to have a high chance of being targeted.

2.3.1 Risk Adjusted Performance

For evaluating the performance of any portfolio we rely on the same methodology as MacKinlay (1997), using a market model (10) to relate the return of the portfolio to that of the market. Portfolio abnormal returns are then computed through equation (11) to (12) and tested statistically in equation (13).

$$\begin{gathered} E(R_{p,t}) = \hat{\alpha}_p + \hat{\beta}_p R_{m,t} + \epsilon_{p,t} \\ E(\epsilon_{p,t}) = 0 \\ Var(\epsilon_{p,t}) = \sigma^2_{\epsilon_t} \end{gathered} \tag{10}$$

where $R_{p,t}$ is expected return for portfolio p during day t, $\hat{\alpha}_p$ and $\hat{\beta}_p$ modelled parameters, and $R_{m,t}$ the return of the market portfolio for the same day, t. $\hat{\alpha}_p$ and $\hat{\beta}_p$ are estimated through a 252-day estimation window the year before the holding period.

$$AR_{p,t} = R_{p,t} - \underbrace{(\hat{\alpha}_p + \hat{\beta}_p R_{m,t})}_{E(R_{p,t})} \tag{11}$$

where $AR_{p,t}$ is the abnormal return for portfolio p during day t, $R_{p,t}$ the actual return for portfolio p during day t.

$$CAR_p(t_1, t_2) = \sum_{t=t_1}^{t_2} AR_{p,t} \tag{12}$$

where $CAR_p(t_1, t_2)$ is simply the sum of abnormal returns for portfolio p during the time window t_1, t_2. We use the broad-market index S&P 1500 Super Composite Index (described in more detail in section 2.1) as a proxy for the market portfolio.

Having obtained cumulative abnormal returns, their respective significance is tested against the null hypothesis of CAR equalling 0 using the test statistic in equation (13).

$$\theta = \frac{CAR_p(t_1, t_2)}{var(CAR_p(t_1, t_2))^{1/2}} \sim N(0, 1) \tag{13}$$

where an estimate of the variance is obtained by using the variance of the daily portfolio abnormal returns during the estimation year and scaling this to obtain yearly variance.

2.3.2 Investment Strategy

We propose a strategy of investing in the beginning of a year, in an equal-weighted portfolio of predicted takeover targets and immediately selling any firms which receive takeover bids on the day of bid announcement. The basis for our proposed strategy lies in previous literature which establish that takeover targets (both successful and unsuccessful) experience large abnormal returns on the day of bid announcement. However, with diverging post announcement performances based on the outcome of the takeover (Bradley et al. (1983)), we propose to sell firms on the day of bid announcement.

To evaluate that our proposed investment strategy carries any initial viability, we will build yearly equally-weighted portfolios containing all target firms from our testing data (2017-2019) and use the methodology in 2.3.1 to compute their risk adjusted performance. While one obviously cannot ex ante identify all takeover targets in the financial markets and compose portfolios of them, the analysis is important to show that our proposed strategy has any merit, disregarding predictive accuracy.

2.3.3 Portfolio Construction

To create investment portfolios for our proposed strategy, an important decision is setting an appropriate cutoff probability for the prediction models. The estimated takeover likelihoods for each firm, produced by the prediction models, will then be compared to the cutoff probability. If the estimated takeover likelihood exceeds the cutoff, the firm will be classified as a future target and included in the portfolio. If not, it will be classified as a future non-target. Previous literature have approached the issue in different ways. Powell (2001) selects the probability cutoff where positive predictive value (i.e. ratio of targets) of the portfolio of predictions is maximized. For this, he uses data from the training datasets. Palepu (1986) instead uses a methodology which aims to minimize misclassification in general, regardless of an observation being target or non-target. The drawback of this methodology is that it does not necessarily focus on maximizing the concentration of actual targets in a portfolio, an essential part in our proposed strategy.

We see an issue in constraining our study to a single cutoff value generated by analysis of the training dataset, as done in previous research. Firstly, we believe the selection of cutoff probabilities should be related to the distribution of takeover likelihoods in the dataset on which predictions are performed rather than on the training dataset (albeit without utilizing information only available ex post - which firms were actually targets or non-targets). Secondly, we see an interesting extension of our study in utilizing a range of cutoffs rather than a single number. For example, a higher probability cutoff should translate into a prediction with higher positive predictive value. On the other hand, increasing the cutoff will most likely have a severe impact on the number of firms included in the portfolio. A range of cutoffs could reveal interesting information about this trade-off and its impact on portfolio returns.

Thus, instead of trying to find an optimal cutoff probability, we utilize a range of different cutoff probabilities. We construct portfolios by looking at the distribution of estimated takeover probability in the test samples from the two prediction models. Four portfolios are constructed with cutoffs at the 80th, 90th, 95th and 99th percentile of the probability distribution of each model. The cutoffs are recalculated each year (2017-2019). In total, this yields 24 cutoff probabilities and 24 portfolios (four cutoffs over three years for two models).

Portfolio cumulative abnormal returns for these portfolios and their respective significance are calculated using the same methodology as proposed by MacKinlay (1997) (shown in section 2.3.1). As we aim to trade on the mechanism of takeover premiums, we further present returns of only the actual takeover targets in the portfolios studied (i.e. our true positives).

Trading portfolios of the test sample are rebalanced annually, meaning the prediction models get one more year of training data for each consecutive year. In other words, 2017's trading portfolio is built by predictions trained by data from 2010 to 2016. The portfolio for 2018 is then rebalanced by predictions trained on data from 2010 to 2017 and so on.

3 Empirical Results

In this section we present the empirical results of our study. We start by analyzing descriptive statistics of target and non-target firms respectively. Next, we present returns and statistical tests of all-target portfolios to prove initial viability of our proposed investment strategy. We then move on to analyze our two prediction models in detail and specifically test if the random forest model brings any additional predictive power. Finally, we present returns and statistical tests for portfolios formed by predictions from the respective models.

3.1 Descriptive Statistics

The purpose of this presentation is both to give the reader an understanding of our dataset and to analyze any differences between target and non-target firms with respect to independent variables.

Table 3 shows a comprehensive overview of descriptive statistics for target and non-target firms over the full sample period, i.e. including both training and test data. Table 13 in appendix further presents a two-tailed t-test, comparing the significance of the differences in means between the two groups. Looking at the tables, target and non-target firms evidently exhibit some different characteristics in independent variables. Firstly, coherent with the inefficient management hypothesis in table 1, underperforming firms seem to be more likely targets of acquisition. Means for Return on equity, net margin, operating margin, and average excess return are all lower for target firms than for non-targets. However, average excess return is the only variable being statistically significant (highly significant), potentially indicating that market based measures do a better job in capturing this effect. The variable turns negative for target firms, while being positive for non-targets, clearly stating a lack in performance for firms being targeted. Sales growth, although insignificant, is pointing in the opposite direction with mean for targets being higher than that for non-target firms.

Both the growth-resource mismatch dummy and the industry dummy variables are highly significant, pointing in the same direction as in previous literature, suggesting that firms with a growth-resource mismatch and firms operating in industries with previously high M&A activity are more likely to receive takeover bids. Size variables (total assets, net sales and market capitalization) are also consistently lower for target firms, indicating that firm size correlates negatively with acquisition likelihood. Out of the three, market capitalization is the only one being statistically significant.

Market-to-book value and price-earnings ratio suggest the opposite compared to previous literature, with that of targets being higher than for non-targets. The dividend yield is the same for both targets and non-targets, but earnings yield is lower for target firms. All these variables are, however, insignificant.

Leverage, both long-term and short-term, is statistically significant and points in the same direction as the literature suggests, with target firms having significantly higher leverage than non-target firms.

Momentum, quite surprisingly, is higher (statistically significant) for non-target firms, differing from previous literature.

Table 3: Descriptive statistics

	Mean		St. Dev.		Min		Max	
	Targets	Non-Targets	Targets	Non-Targets	Targets	Non-Targets	Targets	Non -Targets
(1) Return on equity	0.065	0.098	1.160	2.360	−20.192	−89.590	7.155	20.243
(2) Sales growth	0.200	0.095	2.624	0.745	−0.684	−0.588	66.127	65.571
(3) Momentum*	1.078	1.218	0.460	7.106	0.194	0.064	4.867	706.494
(4) Liquidity	−0.040	−0.054	0.228	0.196	−0.873	−1.487	0.863	0.842
(5) Cash-to-capital ratio	0.549	0.274	4.895	3.671	−10.704	−292.816	95.220	120.554
(6) Long-term leverage*	2.756	0.718	26.036	30.091	−35.607	−1, 601.305	421.733	2, 170.615
(7) Short-term leverage*	1.379	0.625	10.827	8.754	−39.329	−580.400	194.412	311.667
(8) Total assets	10, 442.430	10, 735.800	29, 671.070	28, 260.960	42.915	13.877	531, 864.000	444, 097.000
(9) Total sales	7, 961.600	8, 358.039	20, 458.440	25, 656.240	0.158	0.175	208, 357.000	514, 405.000
(10) Market capitalization***	9, 525.067	12, 856.980	22, 472.700	40, 133.960	39.476	6.546	225, 905.800	1, 073, 391.000
(11) Market-to-book ratio	9.434	3.438	135.209	43.898	−63.487	−2, 182.839	3, 399.450	1, 300.629
(12) Price-earnings ratio	74.257	69.308	247.621	419.909	−1, 215.500	−11, 271.000	3, 342.333	9, 399.000
(13) Dividend yield	0.014	0.014	0.034	0.026	0.000	0.000	0.484	0.913
(14) Earnings yield	−0.012	−0.003	0.170	0.220	−3.358	−14.129	0.590	1.079
(15) Net margin	−1.197	0.006	22.553	3.388	−517.234	−335.029	0.544	3.196
(16) Operating margin	−1.156	0.071	23.128	2.819	−541.481	−278.446	0.664	0.682
(17) Growth-resources dummy**	0.444	0.397	0.497	0.489	0	0	1	1
(18) Industry disturbance dummy***	0.489	0.388	0.500	0.487	0	0	1	1
(19) Average excess return***	−0.00002	0.0001	0.001	0.001	−0.005	−0.010	0.003	0.013
Number of Target Observations	640							
Number of Non-Target Observations	9, 786							
Total Observations	10, 426							

This table presents descriptive statistics for target and non-target firms in our sample. The sample composes firms that were constituents of the S&P 1500 Super Composite Index between 2010-2019, excluding firms in the financial services industry. Using SDC Platinum's classification of M&A, we identify 640 targets and 9,786 non-targets over the period. Note that since the time window is rolling, firms recur as observations multiple times, describing the magnitude of non-target observations. A company is defined as a target if it got a takeover announcement the year following the observation, i.e. explanatory variables are measured end of year the year before announcement. Monetary measures are in millions of US dollars. Variables are defined to following way: *(1) Return on equity:* Four year moving average of a firm's return on equity *(2) Sales growth:* three year moving average of the year-over-year net sales growth *(3) Momentum:* year-over-year change in a firm's stock price *(4) Liquidity:* cash and short-term investments less current liabilities, divided by total assets *(5) Cash-to-capital ratio:* cash relative to shareholders' equity *(6) Long-term leverage:* total long-term debt relative to shareholders' equity *(7) Short-term leverage:* total current liabilities relative to shareholders' equity *(8) Total assets:* total assets in millions of \$ *(9) Total sales:* total net sales in millions of \$ *(10) Market capitalization:* total market capitalization in millions of \$ *(11) Market-to-book ratio:* the ratio of a firm's equity at market value to its equity at book value *(12) Price-earnings ratio:* share price relative to net income *(13) Dividend yield:* dividend relative to share price *(14) Earnings yield:* earnings per share relative to share price *(15) Net margin:* net income divided by net sales *(16) Operating margin:* operating profit divided by net sales *(17) Growth-resources dummy:* dummy variable assuming value 1 if a firms has a "growth-resource mismatch" and 0 otherwise. A "growth-resource mismatch" is defined as having above sample average sales growth but below sample average of resources, or vice versa. Resources are defined as a firm's cash and short-term investments less its total current liabilities, divided by total assets *(18) Industry disturbance dummy:* dummy variable assuming value 1 if there has been an announced takeover in a firm's four-digit SIC code during the year. If no such takeover has been announced, it assumes value 0 *(19) Average excess return:* Four year moving average of daily abnormal returns of a firm's stock. Daily abnormal returns are calculated using a market model where beta is estimated on a 252-day trailing window
(*) denotes significance of means comparison (presented in table 13 in appendix) where *p<0.1; **p<0.05; ***p<0.01

Table 4: Number of target and non-target firms per year

Year	Targets	Non-Targets	Target ratio
2011	54	961	5.62%
2012	64	1,130	5.66%
2013	60	1,133	5.30%
2014	70	1,115	6.28%
2015	87	1,088	8.00%
2016	87	1,098	7.92%
2017	81	1,097	7.38%
2018	71	1,084	6.55%
2019	66	1,080	6.11%
Total	640	9,786	6.54%

This table presents number of target and non-target observations for each consecutive year. The target ratio is computed as the number of targets to total observations the same year.

Table 4 shows the number of target and non-target observations for each year. Not only does this provide an overview of the data, but also highlights one of the main issues in predicting takeover announcements; the data imbalance. For our whole sample, only about 7% of companies receive takeover bids during the period. This puts a special emphasis on the choice of model evaluation metric. For instance, evaluating a model based solely on its accuracy (observations correctly classified divided by total observations) would lead to the selection of a model that predicts 0 targets for the entire period, with an astonishing accuracy of 93%, but useless in the hands of a fund manager.

3.2 Target Portfolio Returns

To prove initial viability of the strategy of investing in takeover targets, we analyze the cumulative abnormal returns of equally weighted portfolios containing only target firms in a given year in table 5 (see table 4 for number of firms for each year). The portfolios mimic our proposed investment strategy, investing in target firms and selling on the day of the bid announcement. The holding period is one year and market model parameters are estimated in the year preceding the holding period. One obviously cannot form portfolios containing only actual takeover targets in a realistic scenario. However, this analysis is important to establish that if one is able to construct a portfolio with a sufficient proportion of true takeover targets, it could indeed constitute an investment strategy earning abnormal returns.

We find that portfolios in all three years 2017-2019 experience positive abnormal returns, all being statistically significant. Combining these findings with the consistent findings in previous research on takeover target returns, we conclude that our proposed investment strategy, together with a powerful target prediction model, could constitute an investment strategy yielding strong abnormal returns.

Table 5: Cumulative abnormal returns all-target portfolios

(%)	2017	2018	2019
CAR	21.22**	58.03***	55.53***

*p<0.10;**p<0.05;***p<0.01

This table presents cumulative abnormal returns (see section 2.3.1 for computation) for our equally-weighted all-target portfolios. The holding period is one year and companies are being sold at the day of the takeover announcement. Their respective significance is tested against the null hypothesis of CAR equalling zero.

3.3 Binomial Logistic Regression

In this section, we present and analyze results from the logistic regression model. We begin with analyzing which independent variables have statistically significant predictive power. Next, we then present and analyze prediction performance measures. Finally, we form portfolios based on the logistic regression model predictions and analyze the compositions of these portfolios.

Table 6 presents the four logistic regression models used by Palepu (1986). Models 1 and 3 use 6 independent variables, but differ as model 1 uses the market-based performance measure average excess return while model 3 relies on accounting data in the form of return on equity. Models 2 and 4 use the same variables as 1 and 3, respectively, but also includes three additional variables; growth, liquidity and leverage. The reason for including these is to showcase that if firms have a growth-resource mismatch, what mismatch is prevalent.

The results of the binomial logistic regression to some degree corresponds to the hypotheses presented by previous literature[4]. Table 6 shows that average excess return and the industry-disturbance dummy are statistically significant and pointing the in the same direction as the literature suggest across all models they are included in. Market-to-book ratio is statistically significant, but proves to have the opposite of the suggested effect on the likelihood of takeover, with firms having a higher ratio also having a higher probability of being acquired.

In model 3 and 4, both the growth-resource dummy and the size variable are statistically significant and coherent with previous literature, i.e. firms with a growth-resource mismatch are more likely to be targeted and takeover likelihood decreases with size. The growth variable itself is also statistically significant and has a positive impact on the takeover likelihood. However, liquidity and leverage are insignificant making it difficult to distinguish what growth-resource mismatch is more frequent.

[4]see table 1 for explanation of these variables and the expected effect on the likelihood of takeover

Table 6: Logistic regression results

	Dependent variable:			
	Company is Targeted			
	(1)	(2)	(3)	(4)
Average excess return	−205.739***	−209.743***		
	(48.657)	(49.003)		
Return on equity			0.0003	0.0005
			(0.018)	(0.018)
Growth-resource dummy	0.136	0.136	0.147*	0.167*
	(0.083)	(0.093)	(0.083)	(0.092)
Growth		0.042**		0.039*
		(0.021)		(0.021)
Liquidity		0.019		−0.097
		(0.233)		(0.230)
Leverage		0.001		0.001
		(0.001)		(0.001)
Industry disturbance dummy	0.414***	0.409***	0.419***	0.421***
	(0.082)	(0.083)	(0.082)	(0.083)
Size (Market capitalization)	−0.040	−0.039	−0.056**	−0.056**
	(0.026)	(0.026)	(0.025)	(0.026)
Market-to-book ratio	0.001**	0.001*	0.001**	0.001**
	(0.001)	(0.001)	(0.001)	(0.001)
Price-earnings ratio	0.0001	0.0001	0.00004	0.00004
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Constant	−2.661***	−2.676***	−2.552***	−2.569***
	(0.213)	(0.213)	(0.212)	(0.213)
Observations	10,426	10,426	10,426	10,426
Log Likelihood	-2,387.016	-2,385.087	-2,395.429	-2,393.697
Akaike Inf. Crit.	4,788.032	4,790.175	4,804.859	4,807.394

*p<0.1; **p<0.05; ***p<0.01

This table is a replication of the work presented by Palepu (1986) utilizing our sample of firms. The sample composes firms that were constituents of the S&P 1500 Super Composite Index between 2010-2019, excluding firms in the financial services industry. Using SDC Platinum's classification of M&A, we identify 640 targets and 9,786 non-targets over the period. Note that since the time window is rolling, firms recur as observations multiple times, describing the magnitude of non-target observations. A company is defined as a target if it got a takeover announcement the year following the observation, i.e. explanatory variables are measured end of year the year before announcement. Variables are defined the following way: *Average excess return:* four year moving average of daily abnormal returns of a firm's stock. Daily abnormal returns are calculated using a market model where parameters are estimated on a 252-day trailing window *Return on equity:* four year moving average of a firm's return on equity *Growth-resources dummy:* dummy variable assuming value 1 if a firms has a "growth-resource mismatch" and 0 otherwise. A "growth-resource mismatch" is defined as having above sample average sales growth but below sample average of resources, or vice versa. Resources are defined as a firm's cash and short-term investments less its total current liabilities, divided by total assets *Growth:* three year moving average of the year-over-year net sales growth *Liquidity:* cash and short-term investments less current liabilities, divided by total assets Leverage: total long-term debt relative to shareholders' equity *Industry disturbance dummy:* dummy variable assuming value 1 if there has been an announced takeover in a firm's four-digit SIC code during the year. If no such takeover has been announced, it assumes value 0 *Size:* total market capitalization in millions of $ *Market-to-book ratio:* the ratio of a firm's equity at market value to its equity at book value Price-earnings ratio: share price relative to net income. Table 14 in appendix further presents a correlation matrix of all these variables. None of the variables exhibit high correlation with one another.

Just as Palepu (1986), model 2 in table 6 is used for the actual prediction tasks. In figure 3 we plot the ROC and present ROC AUC of the logistic regression model for both the out-of-bag proportion of the first training dataset (years 2010-2016) and the three test datasets (years 2017-2019). The ROC AUC for the training dataset is 0.59 and for the test datasets 0.63, 0.53 and 0.66 respectively. This is interpreted as a model with a 63%, 53% and 66% chance respectively to assign a higher takeover probability to an actual target firm than to an actual non-target firm. For 2017 and 2019, the ROC AUCs signify a model which holds some predictive power. For 2018, this is less certain.[5]

Figure 3: ROC logistic regression model

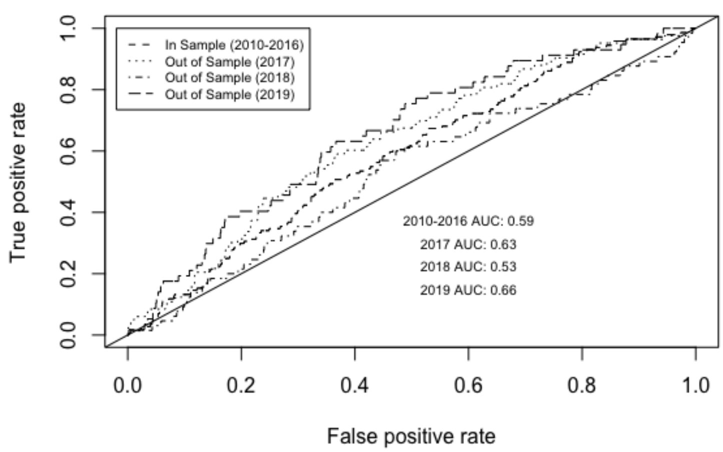

This figure shows the ROC curves (receiver operating characteristic; a plot of true positive rate against false positive rate over many cutoff probabilities) and the ROC AUCs for our logistic regression model. The metrics are shown for all test datasets (2017-2019) and for the out-of-bag portion of the first training dataset (2010-2016). ROC AUC is here interpreted as the chance of the model to assign a higher probability of takeover to an actual target than to an actual non-target. Judging from the ROC AUC values, the model seems to have predictive value in 2010-2016, 2017 and 2019 but maybe not in 2018.

The distributions of targets and non-targets with respect to logistic regression estimated takeover probabilities (figure 4) show quite tight distributions between approximately 0.0 and 0.20. The two groups overlap to a high degree. However, targets evidently have a higher proportion of observations with higher estimated takeover probability, evidenced by the slight right-shift in the target distribution relative to the non-target distribution. This pattern is expected, given a ROC AUC of 0.53 to 0.66 (a ROC AUC of 1 would mean completely separate distributions, while 0.5 would mean completely overlapping distributions). Interestingly, the model is able to assign very high probabilities to a small number of observations which also turn out to be targets (evidenced by few right-tail observations).

[5] a Wilcoxon Rank Sum test proves significance below 1% level for the 2017 and 2019 datasets, proving differences in the distribution of takeover likelihood between target and non-target firms. See table 15.

Figure 4: Logistic regression estimated takeover probability distribution

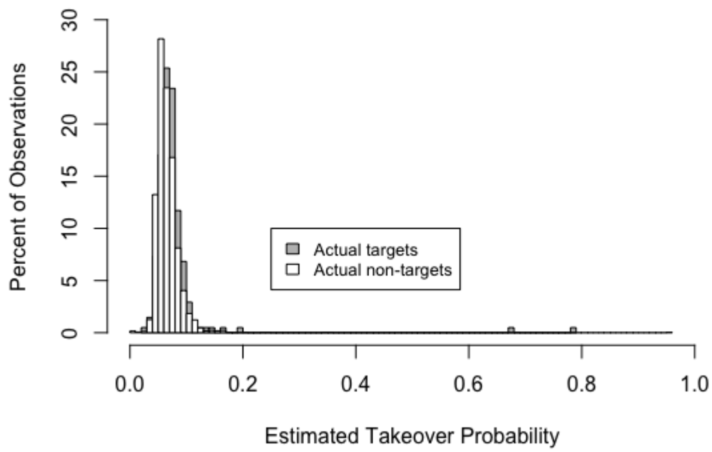

The figure above shows the distribution of estimated takeover probabilities generated by the logistic regression model for our target and non target firms in the combined test datasets (2017-2019). Targets and non-target firms are highly overlapping however the distribution of targets seems to have a slight right shift, indicating the model assigns a slightly higher estimated probability of takeover to this group of firms.

In table 7, we show compositions of the portfolios of firms constructed through the logistic regression prediction model. We construct four portfolios per test dataset, each with a different probability cutoff based on the 80th, 90th, 95th and 99th percentiles of the estimated takeover probability distribution each year. Using the information in table 4 on the percentage of targets per year in the sample as a benchmark, the logistic regression model consistently constructs portfolios with a higher percentage of targets than for the total sample in the corresponding year. The difference is often slim, however, even for higher cutoff probabilities. The model performs best in 2017 and worse in 2018 and 2019. For 2019 this is quite surprising, given the relatively high ROC AUC in this year.

To evaluate which probability cutoff that yields the highest PPV we weigh the PPVs for all three test datasets (2017-2019) and find that the highest percentile, the 99th, performs best with a weighted average PPV of 11.11%. This speaks in favor of the model, as a well-performing prediction model should in general see an increase in positive predictive value as the cutoff probability is increased. Analyzing the effect of cutoff probability on the number of firms in the portfolio, we see that the number of firms decrease drastically with an increased cutoff as expected. In the 99th percentile portfolio, the number of firms is only twelve each year with one to two actual targets. With the weak predictive power of the model, it could be discussed whether choosing a lower probability cutoff (e.g. the 95th percentile portfolio) would be a wiser course of action. We return to this question in section 3.6 when evaluating portfolio returns over the years.

Table 7: Logistic regression predictions 2017-2019

	80^{th} percentile			90^{th} percentile		
	Targets	Non-targets	PPV (%)	Targets	Non-targets	PPV (%)
2017	24	212	10.17	11	107	9.32
2018	19	212	8.23	9	107	7.76
2019	19	211	8.26	9	106	7.83
PPV all years			8.90			8.31

	95^{th} percentile			99^{th} percentile		
	Targets	Non-targets	PPV (%)	Targets	Non-targets	PPV (%)
2017	8	51	13.56	1	11	8.33
2018	4	54	6.90	2	10	16.67
2019	5	53	8.62	1	11	8.33
PPV all years			9.71			11.11

This table presents compositions of portfolios constructed by our logistic regression predictions. Portfolios are built by looking at the distribution of estimated takeover probabilities. Positive predictive value (PPV) is the ratio of targets to total number of firms in the portfolio. See table 16 for exact cutoff probabilities used.

Summarizing our findings, we see that the logistic regression model has some predictive power in the context of predicting future takeover targets. Average excess return and an industry disturbance dummy seem to be the two most important independent variables. The model can be used to consistently construct portfolios with a higher target ratio than the sample itself. A relatively higher cutoff probability is advantageous compared to a lower one with respect to positive predictive value but also drastically decreases the number of firms in the portfolio.

3.4 Random Forest

In this next section, we present the empirical results from the random forest prediction model. The disposition mimics the logistic regression section with additional analysis comparing the two models to one another.

In figure 5 we present the ROC curve and the ROC AUC of the random forest model. The metrics are presented for both the out-of-bag portion of the first training dataset (years 2010-2016) and the three test datasets (years 2017-2019). A ROC AUC of 0.59 for the out-of-bag portion of the training dataset and 0.63, 0.57 and 0.66 respectively for the three test datasets are seen. As with the logistic regression model, this can be interpreted as a model with a 63%, 57% and 66% chance respectively to assign a higher takeover likelihood to an actual target firm than to an actual non-target firm. Compared with the non-discrimination benchmark of 0.50, clearly the random forest holds predictive power in years 2017 and 2019. In 2018, there seems to be some predictive power however this is less certain.[6] Compared to the logistic regression model, the random forest appears to perform very similar for all datasets except for the second test dataset (year 2018) when it outperforms logistic regression. Given the amount of literature suggesting better performance of random forest for imbalanced classification problems, the lack of difference in performance is somewhat surprising.

Figure 5: ROC Random forest model

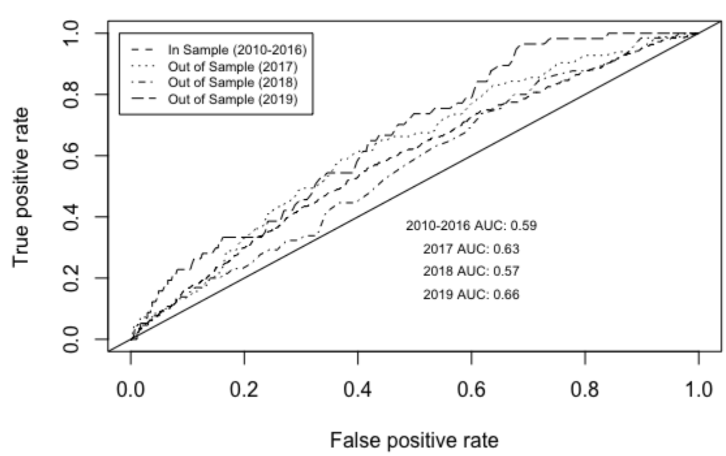

This figure shows the ROC curves (receiver operating characteristic; a plot of true positive rate against false positive rate over many cutoff probabilities) and the ROC AUCs for our random forest model. The metrics are shown for all test datasets (2017-2019) and for the out-of-bag portion of the first training dataset (2010-2016). ROC AUC is here interpreted as the chance of the model to assign a higher probability of takeover to an actual target than to an actual non-target. Judging from the ROC AUC values, the model seems to have some predictive value in all years with the best predictive ability in 2019.

[6]see table 15 for statistical testing of the models' discriminatory abilities.

Figure 6 displays the probability distributions of targets and non-targets with respect to random forest estimated takeover probability. The distributions are based on the combined test datasets. Interestingly, the random forest model produces distributions of estimated takeover likelihood that are much more right-skewed than those of the logistic regression model. This difference in distributional shape corroborates our choice of using percentiles of the model distributions as cutoff probabilities, rather than using some arbitrary value, as the distributions differ so much. For the random forest model, both targets and non-targets exhibit pronounced right skews ("fat tails"), however the distribution of actual targets is shifted to the right compared to that of non-targets. As for the logistic regression model's distributions, this slight right-shift is expected for a model with a small predictive value.

Figure 6: Random forest estimated takeover probability distribution

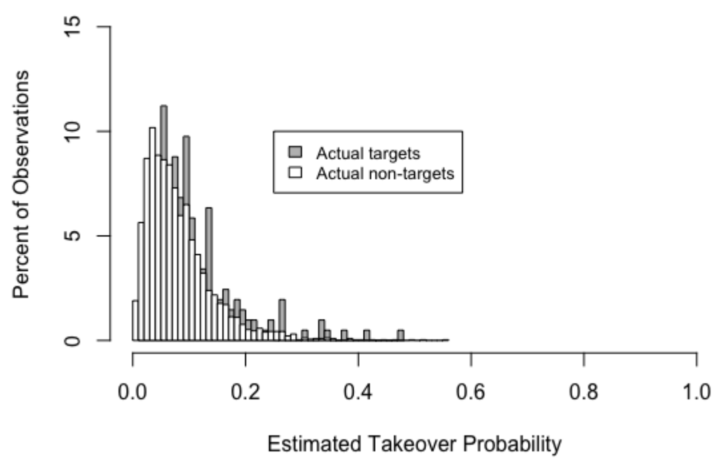

The figure above shows the distribution of estimated takeover probabilities generated by the random forest model for our target and non target firms in the combined test datasets (2017-2019). Evidently, the random forest model produces distributions which are right-skewed. Targets and non-target firms are overlapping but the distribution of targets seems to have a slight right shift, indicating the model assigns a slightly higher estimated probability of takeover to this group of firms. 4.

In figure 7 we plot the relative importance of independent variables used in the random forest model. Many of the variables exhibit similar importance, but momentum, liquidity and average excess return stands out as being especially important. Given the highly significant difference in means between targets and non-targets for average excess return (see table 13), the fact that this variable contributes highly to Gini index is not surprising. More surprising is that momentum and liquidity are ranked as such important variables by the random forest model, given that the mean differences between targets and non-targets are much less significant.

Some variables appear to not contribute much to the predictive model. This includes dividend yield, industry-disturbance dummy and growth-resource-dummy. This is somewhat interesting, especially as the dummy variables exhibit highly significant differences in means between targets and non-target firms. Industry disturbance dummy is also highly significant in the logistic regression model.

Figure 7: Random forest variable importance

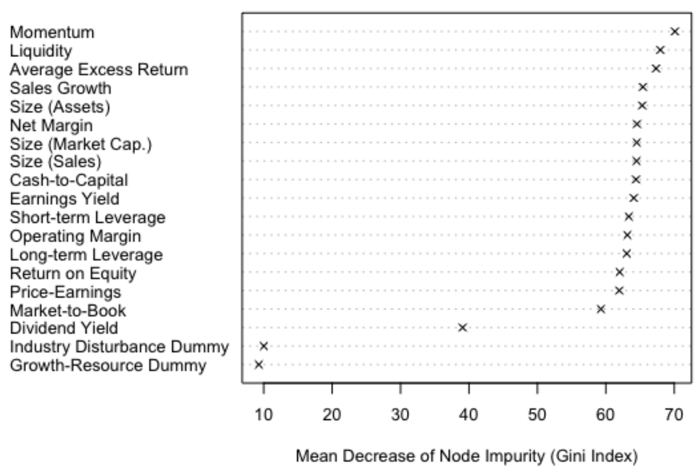

The figure above shows the importance of independent variables used in the random forest model, as measured by their mean decrease of Gini Impurity over all ensemble trees. The data comes from the first training dataset (2010-2016). Momentum, liquidity and average excess return are the most important variables while the two dummy variables exhibit the least importance.

In table 8 we present the compositions of the portfolios constructed through the random forest prediction model. As for the logistic regression model, we create four portfolios for every test dataset. The four portfolios differ in their individual probability cutoffs, which are set according to the percentile of estimated takeover probability distribution in each year (80th, 90th, 95th and 99th percentile). Studying table 8, it is apparent that the random forest produces portfolios with a higher than sample target ratio (PPV) in all cases but one (95th percentile, 2017).

Benchmarking the PPVs of the random forest model to the logistic regression model, the random forest outperforms the latter in six of the portfolios and is at par with the same in two portfolios. In four portfolios, the logistic regression model outperforms the random forest model. The difference in performance varies clearly between years. In 2019, the random forest model outperforms logistic regression for every single cutoff. For 2017, the logistic regression model seem superior instead.

To evaluate which probability cutoff that yields the highest PPV we weigh the PPVs for all three test datasets (2017-2019) and finds that the 99th percentile, the highest, performs best with a weighted average PPV of 13.89%. However, as for the logistic regression model, a higher percentile leads to a drastically decreased number of firms in the portfolio with only twelve firms each year of which one to two are actual targets. Given the weak predictive power of the model, arguably it could be wiser to choose the second best (with respect to PPV) percentile portfolio, i.e. the 90th percentile. This will be further evaluated in section 3.6.

Table 8: Random forest predictions 2017-2019

	80^{th} percentile			90^{th} percentile		
	Targets	Non-targets	PPV (%)	Targets	Non-targets	PPV (%)
2017	23	217	9.58	12	108	10
2018	16	217	6.87	8	108	6.90
2019	23	212	9.79	18	99	15.38
PPV all years			8.76			10.76

	95^{th} percentile			99^{th} percentile		
	Targets	Non-targets	PPV (%)	Targets	Non-targets	PPV (%)
2017	4	57	6.56	1	11	8.33
2018	5	54	8.47	2	10	16.67
2019	8	50	13.79	2	10	16.67
PPV all years			9.55			13.89

This table presents portfolio compositions constructed by our random forest predictions. Portfolios are built by looking at the distribution of estimated takeover probabilities. Positive predictive value (PPV) is the ratio of targets to total number of firms in the portfolio. See table 16 for exact cutoff probabilities used.

Summarizing our findings regarding the random forest model, we see that it brings significant predictive value in all years studied (albeit with less certainty in 2018). Momentum, liquidity and average excess return appear to be the most important independent variables. Using the predictions from the model to form portfolios, one can generally obtain a higher target ratio than for the sample as a whole. The difference between the two models are less than expected, something that is surprising given research pointing out random forest as superior in cases of severe data imbalance.

3.5 Comparison of Models

Despite previous research suggesting a superiority of random forest for imbalanced classification problems, the random forest model and logistic regression ROC AUCs appear to produce quite similar results. For years 2010-2016, 2017 and 2019, the ROC AUCs are equal at two decimals. In 2018 however, random forest clearly performs better. Pooling the test datasets (2017-2019), we find a combined ROC AUC of 0.605 for the logistic regression and 0.617 for the random forest. To test whether the difference in ROC AUCs is statistically significant we employ a two-sided DeLong test. Results are shown in table 9. We find that the difference in performance, measured by the ROC AUC, is not statistically significant.

Table 9: DeLong test

AUC of combined test datasets			
Logistic Regression	Random Forest	Test Statistic	P-Value
0.605	0.617	-0.506	0.613

This table presents the ROC AUC for both models. Test statistic and p-value are from a two-tailed DeLong test, testing the significance in difference of ROC AUC between the two models.

3.6 Investment Strategies

Using the portfolio compositions in table 7 for the logistic regression model and table 8 for the random forest model, we proceed to calculate CARs of these portfolios and subsequently test these returns for significance using the null hypothesis of CAR equalling zero. The calculation and testing of CAR follows our methodology in section 2.3.1.

CARs of the portfolios constructed through the logistic regression model are shown in table 10. These CARs are also decomposed into portfolios containing only actual targets. The logistic regression model creates five portfolios which generate significant positive CAR (three below the 1% level and two below the 5% level). In four of these portfolios, positive CAR seems to be driven by actual targets in the portfolios which exhibit impressive, significant CAR. The fact that most of the portfolios that exhibit significant CAR are in 2018 implies less stability over time and thus a less robust model.

Table 10: Logistic regression portfolios

	80^{th} percentile		90^{th} percentile		95^{th} percentile		99^{th} percentile	
(%)	Targets	Total	Targets	Total	Targets	Total	Targets	Total
2017	14.92	−16.31	29.31	−20.32	−6.10	−22.47	21.80	13.30
2018	129.98***	33.14***	121.92***	46.54**	145.48***	66.23**	136.90**	64.17**
2019	38.47**	10.22	19.25	15.45	3.21	17.28	−10.87	72.72***

This table shows cumulative abnormal returns (see section 2.3.1 for computation) for equally-weighted portfolios built using predictions from the logistic regression model. Furthermore, portfolios containing only the actual takeover targets of each respective portfolio are presented. The four different probability cutoffs are set by looking at the distribution of estimated takeover probabilities. The portfolio holding period is one year and target firms are being dropped on the day of announcement. See table 16 for exact cutoff probabilities used.

Previous literature have discussed the issue with prediction models predicting more obvious targets which experience lesser returns on bid announcement day as bids are anticipated by financial markets. Comparing the CAR of the actual targets predicted by the logistic regression model to the all-target portfolios in table 5, we conclude that this does not seem to apply to our application of the logistic regression model. Instead, this model constructs portfolios with targets which in several cases exhibit stronger CAR than the all-target portfolios.

In section 6, we hypothesized that the highest percentile portfolio may be a suboptimal choice due to its small number of firms, putting special emphasis on the announcement day returns of those few firms. Looking at the returns of actual targets only, we find some support for this hypothesis in that both the 95th and 99th percentile portfolios have years where targets actually generate negative abnormal returns. The lower percentile portfolios (80th and 90th) do not. However, few of the target returns are statistically significant and thus it is hard to draw any certain conclusions from this.

An equivalent presentation for the random forest model portfolios is shown in table 11. These CARs are also decomposed into portfolios containing only actual targets. For the random forest model, we observe only two portfolios with significant CAR driven by actual targets (and only one of these below the 1% level). Both of these are in 2018. However, we also observe four portfolios with significantly negative CAR. In these respective portfolios, abnormal returns of actual targets are not statistically significant different from 0.

Table 11: Random forest portfolios

	80^{th} percentile		90^{th} percentile		95^{th} percentile		99^{th} percentile	
(%)	Targets	Total	Targets	Total	Targets	Total	Targets	Total
2017	11.20	−27.30**	7.29	−26.99*	9.44	−28.61*	15.26	−40.08**
2018	57.95***	17.85	92.17***	22.61	127.67**	32.05*	93.15***	59.12***
2019	71.83***	12.26	60.64***	19.76	8.57	14.07	−14.63	−21.27

This table shows cumulative abnormal returns (see section 2.3.1 for computation) for equally-weighted portfolios built using predictions from the random forest model. Furthermore, portfolios containing only the actual takeover targets of each respective portfolio are presented. The probability cutoff is set by looking at the distribution of estimated takeover probabilities. The portfolio holding period is one year and target firms are being dropped on the day of announcement. See table 16 for exact cutoff probabilities used.

Returning to the discussion on optimal cutoff probabilities, we see a similar pattern as for logistic regression. The highest percentile portfolios leads to smaller portfolios with the risk of the actual targets included in the portfolio not being the ones experiencing the best announcement day returns. This is evidenced by the two lower percentile portfolios (80th and 90th) having four significant all-target portfolios while the two higher percentile portfolios (95th and 99th) only have two and the 99th percentile portfolio even seeing a year when actual targets produce negative abnormal returns (albeit not statistically significant).

Benchmarking the portfolios constructed through the random forest model to the equivalent logistic regression portfolios, the former do not seem to yield favorable returns. Rather, the logistic regression model is more robust across different cutoff probabilities with all four portfolios in 2018 generating significant positive CAR driven by actual targets visa-vi only two for the random forest model portfolios. The logistic regression model also outperforms the random forest model in 2019 where the logistic regression model achieves one statistically significant CAR-positive portfolio and the random forest achieves none.

Summarizing our findings from the portfolio evaluation, we conclude that a minority of the portfolios formed do see statistically significant abnormal returns. The contribution of actual target firms to these returns appear to perhaps be more stable at lower cutoff probabilities, due to the higher number of firms and thus actual targets included in the portfolios, however the statistical significance of this is small. The random forest model does not seem to yield returns favorable to that of the logistic regression model.

4 Conclusion and Further Remarks

4.1 Conclusion

In this study, our aim has been to evaluate whether a successful investment strategy can be founded in takeover target prediction. Previous research has produced mixed results, however the most influential paper in the area rejected the idea that such a strategy could produce significant abnormal returns. Our addition to this body of research lies in including more predictive variables and using statistical techniques, especially the random forest algorithm, which were not available around the time of publication of some earlier papers. Furthermore, we utilize a longer time window of three years (compared to previous studies looking at single years) and study a range of different cutoff probabilities based on model probability distributions instead of restricting ourselves to a single cutoff value.

We start with evaluating whether forming all-target portfolios and holding these firms until the day of bid announcement produces abnormal returns. We find that it does indeed yield impressive abnormal returns, with a high degree of statistical certainty. Combining this with the large body of research in the area of takeover announcement returns, we conclude that the strategy of takeover target prediction could be viable given a good enough prediction model.

Concerning prediction models, we conclude that both the logistic regression model and the random forest model have some predictive power. We base this conclusion on the reported ROC AUCs and the significant Wilcoxon Rank-Sum tests. However, the predictive power is small. The highest ROC AUC achieved is a mere 0.66, interpreted as only a 66% chance for the models to assign a higher takeover likelihood to an actual target firm than to an actual non-target firm. Regarding the comparison of the two models, we find that the difference in ROC AUC is not statistically significant. We conclude that the addition of new independent variables and the use of the random forest algorithm can not be said to produce a significantly better prediction model.

Regarding the possibility of earning abnormal returns from takeover target prediction, we are less cautious than previous literature in rejecting the viability of this strategy. Extending the time window studied and utilizing a range of cutoff probabilities, we find that our logistic regression model produces significant abnormal returns in a third of portfolios studied. The random forest algorithm also produces a few portfolios with significant abnormal returns, although not as many as the logistic regression model. Our conclusion is that our addition of new independent variables and a new statistical technique does not significantly improve the possibility of earning abnormal returns from takeover target prediction. Instead, the standard approach of using logistic regression yields as good or better results.

In earlier research, the two main explanations behind why takeover target prediction does not produce significant abnormal returns are lack of a high enough ratio of actual targets in the predicted portfolios and that prediction models tend to predict more obvious targets whose stock prices thus react less on announcement day. We find that the first explanation may be correct but warrants further research. We however find no empirical support for the second explanation. In the four portfolios (constructed through logistic regression) earning significant abnormal returns, actual targets seem to indeed drive these returns. Oftentimes, the abnormal returns of predicted targets are higher than for the all-target portfolios, further questioning the explanation that predicted targets should be more obvious and more anticipated by financial markets.

4.2 Improvements and Further Research

Given the inconclusive findings regarding the possibility to earn abnormal returns, we see a value in expanding the body of research in this area. First and foremost, it would be interesting to utilize an even longer time window to obtain more certainty in findings and also discover any trends over time. Second, a valuable addition would be to look at how the predictive power of the model varies during the year within the same test data, i.e. does the accuracy of predictions change as time progresses within the test window? These findings could potentially help in selecting a more optimal strategy for when to rebalance portfolios and trade on new predictions.

Regarding prediction models in themselves, there is ample opportunity to attempt to improve models. Newer machine learning approaches, such as popular gradient boosting frameworks, could be attempted. An attempt to find more independent variables, with a focus on predictive value rather than relying on hypotheses surrounding the supposed explanatory value of variables, could also be attempted.

References

Andrade, G., Mitchell, M., and Stafford, E. (2001). New evidence and perspectives on mergers. *Journal of Economic Perspectives*, 15(2):103–120.

Betton, S., Eckbo, B. E., and Thorburn, K. S. (2008). Corporate takeovers. In *Handbook of Empirical Corporate Finance*, pages 291–429. Elsevier.

Bradley, A. P. (1997). The use of the area under the roc curve in the evaluation of machine learning algorithms. *Pattern Recognition*, 30(7):1145–1159.

Bradley, M., Desai, A., Kim, E. H., et al. (1983). The rationale behind interfirm tender offers: Information or synergy? *Journal of Financial Economics*, 11(1-4):183–206.

Brar, G., Giamouridis, D., and Liodakis, M. (2009). Predicting european takeover targets. *European Financial Management*, 15(2):430–450.

Breiman, L. (2001). Random forests. *Machine Learning*, 45:5–32.

Breiman, L., Friedman, J., Stone, C. J., and Olshen, R. A. (1984). *Classification and regression trees*. CRC press.

Brown, I. and Mues, C. (2012). An experimental comparison of classification algorithms for imbalanced credit scoring data sets. *Expert Systems with Applications*, 39(3):3446–3453.

Couronné, R., Probst, P., and Boulesteix, A.-L. (2018). Random forest versus logistic regression: a large-scale benchmark experiment. *BMC Bioinformatics*, 19(1):270.

DeLong, E. R., DeLong, D. M., and Clarke-Pearson, D. L. (1988). Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach. *Biometrics*, pages 837–845.

Goergen, M. and Renneboog, L. (2004). Shareholder wealth effects of european domestic and cross-border takeover bids. *European Financial Management*, 10(1):9–45.

Jarrell, G. A. and Poulsen, A. B. (1989). The returns to acquiring firms in tender offers: Evidence from three decades. *Financial Management*, pages 12–19.

Langetieg, T. C. (1978). An application of a three-factor performance index to measure stockholder gains from merger. *Journal of Financial Economics*, 6(4):365–383.

MacKinlay, A. C. (1997). Event studies in economics and finance. *Journal of Economic Literature*, 35(1):13–39.

Mandelker, G. (1974). Risk and return: The case of merging firms. *Journal of Financial Economics*, 1(4):303–335.

Muchlinski, D., Siroky, D., He, J., and Kocher, M. (2016). Comparing random forest with logistic regression for predicting class-imbalanced civil war onset data. *Political Analysis*, 24(1):87–103.

Palepu, K. G. (1986). Predicting takeover targets: A methodological and empirical analysis. *Journal of Accounting and Economics*, 8(1):3–35.

Powell, R. G. (2001). Takeover prediction and portfolio performance: A note. *Journal of Business Finance & Accounting*, 28(7-8):993–1011.

Ruback, R. S. and Jensen, M. C. (1983). The market for corporate control: The scientific evidence. *Journal of Financial Economics*, 11:5–50.

Schwert, G. W. (1996). Markup pricing in mergers and acquisitions. *Journal of Financial Economics*, 41(2):153–192.

5 Appendix

Table 12: Computation independent variables

(1) Return on equity:	Four year moving average of a firm's return on equity
(2) Sales growth:	Three year moving average of the year-over-year net sales growth
(3) Momentum:	Year-over-year change in a firm's stock price
(4) Liquidity:	Cash and short-term investments less current liabilities, divided by total assets
(5) Cash-to-capital ratio:	Cash relative to shareholders' equity
(6) Long-term leverage:	Total long-term debt relative to shareholders' equity
(7) Short-term leverage:	Total current liabilities relative to shareholders' equity
(8) Total assets:	Total assets in millions of $
(9) Total sales:	Total net sales in millions of $
(10) Market capitalization:	Total market capitalization in millions of $
(11) Market-to-book ratio:	The ratio of a firm's equity at market value to its equity at book value
(12) Price-earnings ratio:	Share price relative to net income
(13) Dividend yield:	Dividend relative to share price
(14) Earnings yield:	Earnings per share relative to share price
(15) Net margin:	Net income divided by net sales
(16) Operating margin:	Operating profit divided by net sales
(17) Growth-resources dummy:	Dummy variable assuming value 1 if a firms has a "growth-resource mismatch" and 0 otherwise. A "growth-resource mismatch" is defined as having above sample average sales growth but below sample average of resources, or vice versa. Resources are defined as a firm's cash and short-term investments less its total current liabilities, divided by total assets
(18) Industry disturbance dummy:	Dummy variable assuming value 1 if there has been an announced takeover in a firm's four-digit SIC code during the year. If no such takeover has been announced, it assumes value 0
(19) Average excess return:	Four year moving average of daily abnormal returns of a firm's stock. Daily abnormal returns are calculated using a market model where parameters are estimated on a 252-day trailing window

This table presents computation of independent variables used in our study. The computation follows the methodology utilized by Palepu (1986) and Brar et al. (2009). Data used for the computations is gathered from the Wharton Research Data Services and the M&A database SDC Platinum.

Table 13: Means comparison target and non-target firms

	Targets	Non-targets	Statistic	P-Value
(1) Return on equity	0.065	0.098	0.630	0.529
(2) Sales growth	0.200	0.095	−1.004	0.316
(3) Momentum	1.078	1.218	1.894	0.058*
(4) Liquidity	−0.040	−0.054	−1.515	0.130
(5) Cash-to-capital ratio	0.549	0.274	−1.399	0.162
(6) Long-term leverage	2.756	0.718	−1.900	0.058*
(7) Short-term leverage	1.379	0.625	−1.727	0.085*
(8) Total assets	10,442.430	10,735.800	0.243	0.808
(9) Total sales	7,961.600	8,358.039	0.467	0.641
(10) Market capitalization	9,525.067	12,856.980	3.416	0.001***
(11) Market-to-book ratio	9.434	3.438	−1.118	0.264
(12) Price-earnings ratio	74.257	69.308	−0.464	0.642
(13) Dividend yield	0.014	0.014	0.289	0.773
(14) Earnings yield	−0.012	−0.003	1.278	0.202
(15) Net margin	−1.197	0.006	1.349	0.178
(16) Operating margin	−1.156	0.071	1.342	0.180
(17) Growth-resources dummy	0.444	0.397	−2.307	0.020**
(18) Industry disturbance dummy	0.489	0.388	−4.983	0.000***
(19) Average excess return	−0.00002	0.0001	4.351	0.000***

*p<0.1; **p<0.05; ***p<0.01

This table presents comparison of means for target and non-target firms in our sample. The sample composes firms that were constituents of the S&P 1500 Super Composite Index between 2010-2019, excluding firms in the financial services industry. Using SDC Platinum's classification of M&A, we identify 640 targets and 9,786 non-targets over the period. Note that since the time window is rolling, firms recur as observations multiple times, describing the magnitude of non-target observations. A company is defined as a target if it got a takeover announcement the year following the observation, i.e. explanatory variables are measured end of year the year before announcement. Monetary measures are in millions of US dollars. Variables are defined to following way: *(1) Return on Equity:* Four year moving average of a firm's return on equity *(2) Sales Growth:* three year moving average of the year-over-year net sales growth *(3) Momentum:* year-over-year change in a firm's stock price *(4) Liquidity:* cash and short-term investments less current liabilities, divided by total assets *(5) Cash-to-Capital ratio:* cash relative to shareholders' equity *(6) Long-term Leverage:* total long-term debt relative to shareholders' equity *(7) Short-term Leverage:* total current liabilities relative to shareholders' equity *(8) Total Assets:* total assets in millions of $ *(9) Total Sales:* total net sales in millions of $ *(10) Market Capitalization:* total market capitalization in millions of $ *(11) Market-to-book ratio:* the ratio of a firm's equity at market value to its equity at book value *(12) Price-earnings ratio:* share price relative to net income *(13) Dividend Yield:* dividend relative to share price *(14) Earnings Yield:* earnings per share relative to share price *(15) Net Margin:* net income divided by net sales *(16) Operating Margin:* operating profit divided by net sales *(17) Growth-resources dummy:* dummy variable assuming value 1 if a firms has a "growth-resource mismatch" and 0 otherwise. A "growth-resource mismatch" is defined as having above sample average sales growth but below sample average of resources, or vice versa. Resources are defined as a firm's cash and short-term investments less its total current liabilities, divided by total assets *(18) Industry disturbance dummy:* dummy variable assuming value 1 if there has been an announced takeover in a firm's four-digit SIC code during the year. If no such takeover has been announced, it assumes value 0 *(19) Average Excess Return:* Four year moving average of daily abnormal returns of a firm's stock. Daily abnormal returns are calculated using a market model where parameters are estimated using a 252-day trailing window

Table 14: Correlation matrix independent variables

	(1) Average excess return	(2) Return on equity	(3) Growth-resource dummy	(4) Growth	(5) Liquidity	(6) Leverage	(7) Industry-disturbance dummy	(8) Size	(9) Market-to-book ratio	(10) Price-earnings ratio
(1) Average excess return	1	0.02	−0.05	0.06	0.07	−0.01	−0.01	0.12	0.02	0.05
(2) Return on equity	0.02	1	−0.02	−0.01	−0.02	0.02	−0.01	0.05	−0.01	0.01
(3) Growth-resource dummy	−0.05	−0.02	1	−0.01	0.46	−0.005	0.07	−0.10	−0.01	0.0003
(4) Growth	0.06	−0.01	−0.01	1	0.08	−0.0004	0.03	−0.01	0.001	−0.003
(5) Liquidity	0.07	−0.02	0.46	0.08	1	−0.003	0.17	−0.14	−0.003	−0.003
(6) Leverage	−0.01	0.02	−0.005	−0.0004	−0.003	1	−0.01	0.003	0.50	−0.01
(7) Industry-disturbance dummy	−0.01	−0.01	0.07	0.03	0.17	−0.01	1	0.06	−0.01	−0.01
(8) Size	0.12	0.05	−0.10	−0.01	−0.14	0.003	0.06	1	0.04	0.03
(9) Market-to-book ratio	0.02	−0.01	−0.01	0.001	−0.003	0.50	−0.01	0.04	1	−0.01
(10) Price-earnings ratio	0.05	0.01	0.0003	−0.003	−0.003	−0.01	−0.01	0.03	−0.01	1

The table above shows a correlation matrix for all the independent variables used for regression in table 6. All data is taken from WRDS. None of the variables exhibit high correlation with one another. Variables are defined the following way: (1) Average excess return: four year moving average of daily abnormal returns of a firm's stock. Daily abnormal returns are calculated using a market model where beta is estimated on a 252-day trailing window (2) Return on equity: four year moving average of a firm's return on equity (3) Growth-resources dummy: dummy variable assuming value 1 if a firms has a "growth-resource mismatch" and 0 otherwise. A "growth-resource mismatch" is defined as having above sample average sales growth but below sample average of resources, or vice versa. Resources are defined as a firm's cash and short-term investments less its total current liabilities, divided by total assets (4) Growth: three year moving average of the year-over-year net sales growth (5) Liquidity: cash and short-term investments less current liabilities, divided by total assets (6) Leverage: total long-term debt relative to shareholders' equity (7) Industry disturbance dummy: dummy variable assuming value 1 if there has been an announced takeover in a firm's four-digit SIC code during the year. If no such takeover has been announced, it assumes value 0 (8) Size: total market capitalization in millions of $ (9) Market-to-book ratio: the ratio of a firm's equity at market value to its equity at book value (10) Price-earnings ratio: share price relative to net income

Table 15: Wilcoxon rank sum test

	Logistic Regression	Random Forest
2017	0.0001	0.0001
2018	0.4130	0.0710
2019	0.0001	0.0000

This table reports p-values from a Wilcoxon Rank Sum test of the estimated takeover probabilities for targets and non-target firms. The test uses a null hypothesis of target and non-target firms belonging to the same distribution regarding estimated takeover probability. Rejecting the null hypothesis proves predictive power above chance for the logistic regression model and random forest model respectively.

Table 16: Cutoff probabilities

	80^{th} percentile		90^{th} percentile	
(%)	Logistic Regression	Random Forest	Logistic Regression	Random Forest
2017	8.18	12.00	9.16	16.50
2018	8.15	12.67	9.68	16.20
2019	7.34	11.30	8.61	14.30

	95^{th} percentile		99^{th} percentile	
(%)	Logistic Regression	Random Forest	Logistic Regression	Random Forest
2017	10.21	20.33	12.86	28.91
2018	11.17	22.00	16.56	34.64
2019	9.61	17.95	12.36	26.96

This table presents cutoff probabilities used in our prediction models. The cutoffs are set by looking at the distributions of estimated takeover likelihood in years 2017-2019 respectively and selecting the values that corresponds to the 80th, 90th, 95th and 99th percentiles.

Figure 8: Number of ensemble trees random forest model

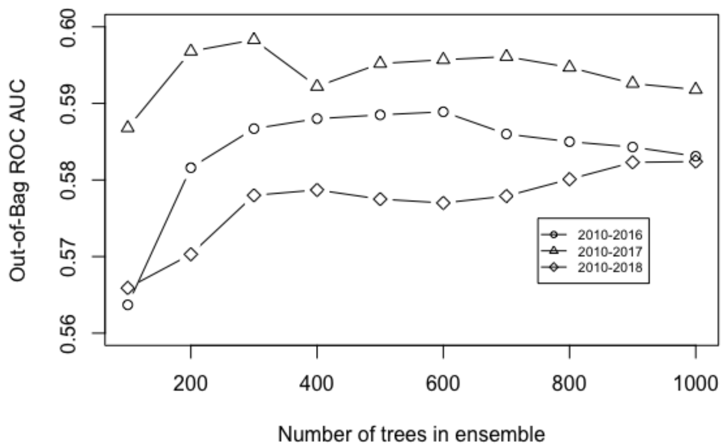

The figure above shows results from tuning the random forest model with respect to number of ensemble trees grown. The mechanism of ensemble trees are described in section 2.2.2. Tuning is done by incrementally increasing the number of trees, optimizing for ROC AUC. We decide to grow 600 ensemble trees for the random forest model trained on data 2010-2016, 300 for the one trained on data from 2010-2017 and 1000 for the one trained on data from 2010-2018.

Figure 9: Number of independent variables tried at each node random forest model

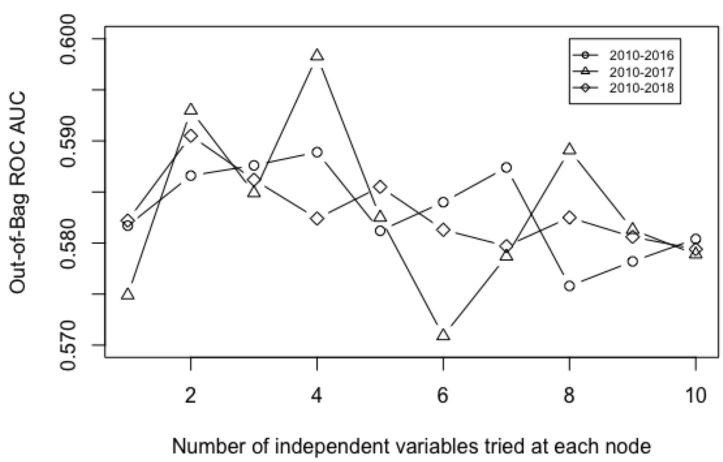

The figure above shows results from tuning the random forest model with respect to the number of independent variables evaluated in each tree node. This parameter is described in more detail in section 2.2.2. Tuning is done by incrementally increasing the number of variables, optimizing for ROC AUC. We decide to set this parameter to four in the random forest model trained on data 2010-2016, four for the one trained on data from 2010-2017 and two for the one trained on data from 2010-2018.