Credit Card Approval and Loyalty Prediction Analysis

James Koh, Kevin Lin, Joshua Lowe, and Aditya Senthilvel

February 21, 2024

Abstract

We propose a four-step process for banks to use to assess customers' financial behavior by analyzing (1) late payments, (2) credit scores, (3) credit card approvals, and (4) credit card churning. Our objective is to predict both credit card approval and customer value. We selected these steps based on several prior beliefs: late payments reflect repayment ability, credit scores influence approval chances, and transaction history predicts customer loyalty. Based on our results, there is a minor negative impact from late payments on credit score, with credit score and transaction history emerging as key factors for credit card approval and customer loyalty. Our predictions achieved cross-validation accuracies of 0.65, 0.68, 0.72, and 0.93 for late payment, credit score, credit card approval, and customer loyalty prediction respectively, beating all baseline metrics of 0.58, 0.53, 0.55, and 0.84 respectively.

1. Introduction or Background

Over 82% of U.S. adults have a credit card, with over half owning two or more. It's widely advised by finance professionals and credit card users to use credit cards responsibly to build a positive credit history. Credit scores play a crucial role in approving applicants for loans and determining interest rates. A higher credit score leads to lower interest rates and better credit card rewards. Many individuals are interested in pre-approval for new credit cards to avoid potential damage to their credit score. Banks issuing credit cards are concerned about users' ability to repay balances and interested in maintaining a long-term relationship, as they profit from interest on unpaid balances.

2. Data Preparation

We obtained four separate datasets that can be tied together by connecting the output of one dataset to the input of another dataset leading us to a clear coherent "story." Our datasets were found on Kaggle or HuggingFace.

For the Payment Lateness dataset, the core inputs were the age of credit and age of the applicant in predicting whether an applicant will be late in paying their credit balance.

For the Credit Score dataset, the core inputs were Outstanding Debt, Annual Income, Credit Mix, Payment Lateness, Payment Behavior, Age of Credit, and Minimum Payment Paid per Month. The output of the inputs is the Credit Score.

For the Credit Approval dataset, the core inputs were income, credit score, and years employed, and the output was whether an applicant was approved for credit.

For the Bank Churners dataset, core inputs were total number of transactions per year, total revolving balance, number of products with the bank, and total transaction amount per year. This is used to predict customer attrition.

The following general preprocessing steps were performed for all datasets:

1. Drop irrelevant columns
2. Drop missing values
3. Drop duplicated values
4. Rename columns to a consistent format
5. Convert columns to appropriate types (e.g. Categorical, Numerical, etc.)
6. Fixed data values with invalid characters that rendered it unusable

A more detailed outline of our data preprocessing follows:

Late Payments Dataset:

1. Merged two datasets on applicant ID: one containing applicant ID, age of credit, and payment lateness; and the other containing personal information for the applicant such as work occupation type, income, housing type, family size, etc. Reduced dataset from 438557 to 30322 records.
2. Binned 'Income' into 5 groups to convert it to be categorical and merged certain 'Highest Edu', 'Housing Type', and 'Income Type' categories (i.e. 'Academic degree' and 'Secondary')
3. Classify 'Payment Lateness' as 'Late' or 'On Time' based on whether it is negative or positive.
4. Dropping more NaN data brings the data to 25134 records of 17 features, excluding applicant ID and the target.

Credit Score Dataset:

1. Dropping all NaN values resulted in the number of records dropping from ~98,000 to 48,000 records.
2. Dropping those values led to 48068 records of 19 features excluding the target.

Credit Approval Dataset:

1. Dropped PriorDefault - not present in other datasets
2. Dropped ZipCode - shouldn't have any effect
3. No NaN/unknown data to remove (no records dropped)
4. 690 records of 13 features, excluding the target.

Bank Churners Dataset:

1. Replaced education and income to be consistent with Late Payments Dataset (i.e. 'More than 120k', and 'Graduate')
2. Dropping unknown data led to a decrease from 101127 records to 7081 records of 19 features, excluding the target.

3. Data Processing

For processing the data, we split the data using a 90-10 ratio for training and testing. For each split, we standardized the numerical-type columns, and one-hot encoded the non-numerical ones. We took 4 approaches: DecisionTreeClassifier, L1 Penalty Logistic Regression, L2 Penalty Logistic Regression, and Elastic Net Regression (combination of L1 and L2). For the DecisionTreeClassifier, we used a cross validated trainer with GridSearchCV, using a hyperparameter for depths from 1 to 10. For the regressions, we tested with 10, 100 and 1000 Cs, max iteration of 1000 (or 10000 if it failed to converge), and saga as the solver. Our baseline metric is using a simple bias regressor, predicting the majority value in the output.

4. Results

Results on Payment Lateness:

Graph 1: Test and Baseline Accuracy of best run for Payment Lateness

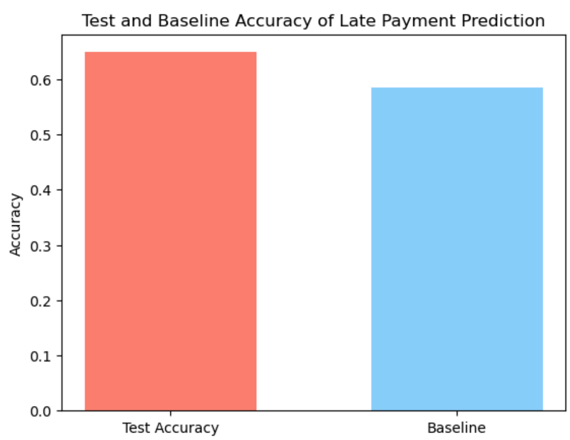

Figure 1: L1 Penalty Logistic Regression Weights for Payment Lateness

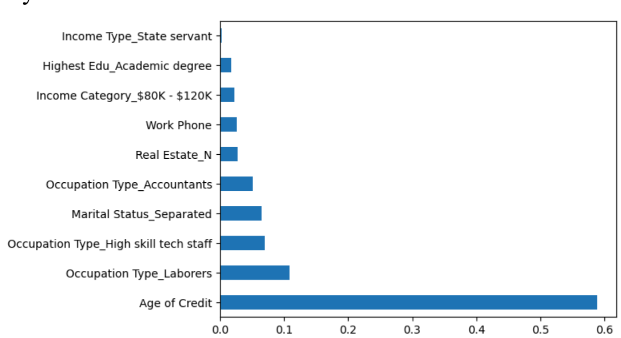

Results on Credit Score:

Graph 2: Test and Baseline Accuracy of best run for Credit Score

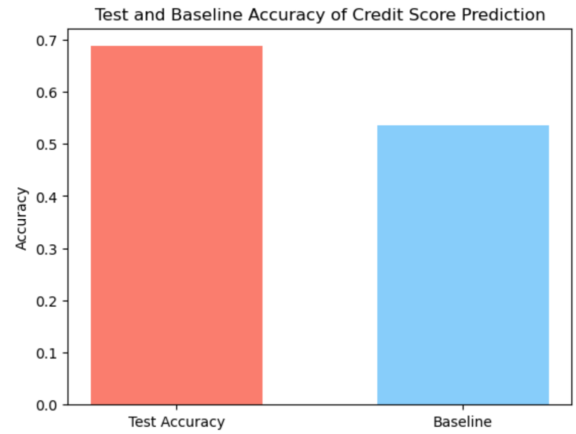

Figure 2: Feature Importance for DecisionTreeClassifer

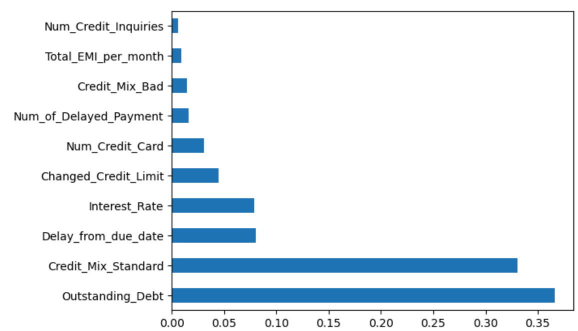

Results on Credit Approval:

Graph 3: Test and Baseline Accuracy of best run for Credit Approval

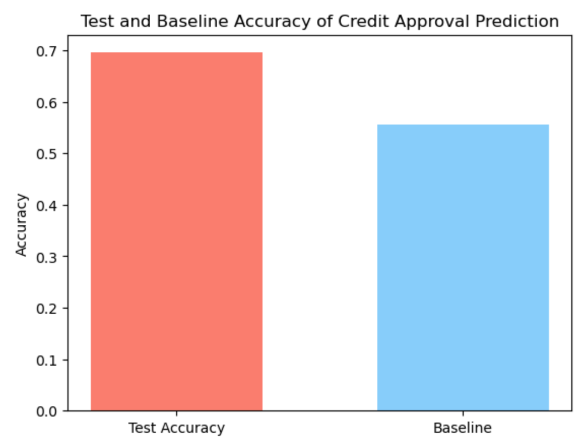

Figure 3: ElasticNet Penalty Logistic Regression Weight for Credit Approval

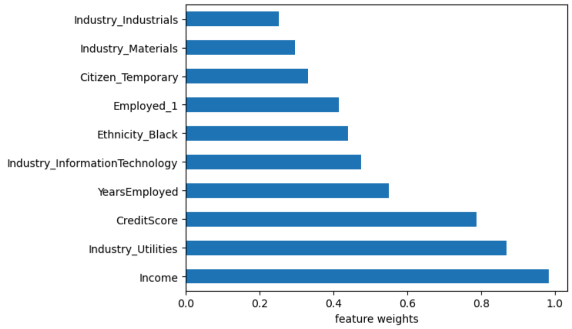

Results on Bank Churners:

Graph 4: Test and Baseline Accuracy of best run for Bank Churners

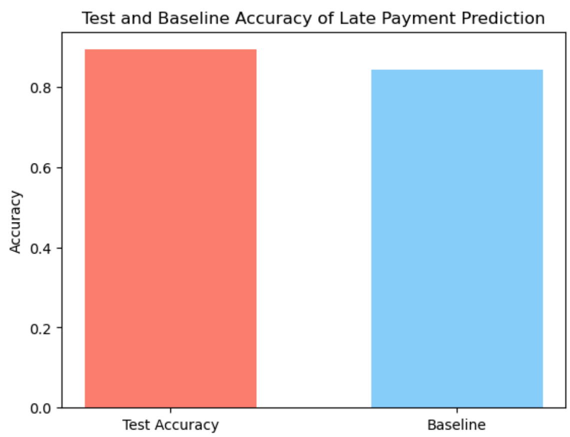

Figure 4: DecisionTreeClassifier Weights for Bank Churners

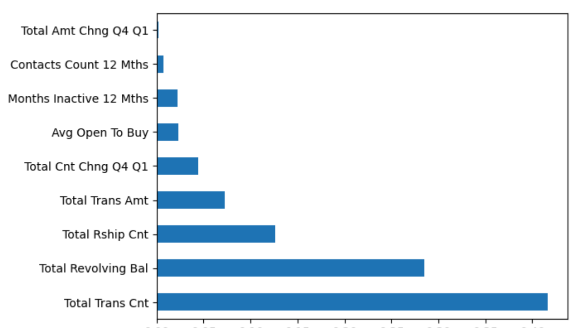

5. Evaluation

For our results on payment lateness, in all processing runs, the age of credit held the most weight and impact on whether a credit card user would likely be late on their payments or not. Depending on the regression run, the decision tree would have the better test accuracy and in other runs the L1 penalty logistic regression run would have the better test accuracy. In either case, the average test accuracy score was around 0.65 which is 0.07 points higher than the baseline of 0.58 as seen in graph 3, surpassing the baseline by roughly 11%.

Analyzing credit score, the dominating factor, as shown in Figure 2, is outstanding debt when determining what the credit score is. While outstanding debt is shown to have an overwhelming importance , payment lateness appears to have an effect as shown in Graph 2. However, since outstanding debt has a much higher importance on credit score, it is possible that while payment lateness can affect credit score, debt will be more likely to have the highest impact. This analysis comes from the fact that the analysis for Figure 2 is approximately 62% accurate, as shown in Graph 2. This does surpass the baseline accuracy by around 10%.

For our results on credit approval, all processing runs resulted in income, credit score, and number of years employed being significant predictors for credit approval. We also found a positive correlation between having a prior default and being approved for a credit card; however, because this data was not present in the other sets, we decided to drop that column. All of the test accuracies were the same between the different penalty metrics, at 0.724, which was about 0.17 points higher than the baseline of 0.55. This is about a 30%

increase in accuracy from the baseline. Since we found previously that credit score was primarily tied to outstanding debt, this means that people who have high income and low outstanding debt while minimizing payment lateness are most likely to be approved for a credit card.

For the results on customers that churn their credit cards, the L1, L2, and ElasticNet regressions reached similar levels of prediction accuracy when using the test split. The decision tree had the highest prediction accuracy of 0.93 and surpassed the others by 0.05. Using the decision tree weights to determine the most important features, customers approved for credit cards who have a high transaction count and low revolving balance are more likely to stay with their bank.

6. Conclusion

We found that the lateness of a payment negatively, alongside outstanding debt affects the credit score of a user and credit score has a significant impact on the outcome of a credit card application getting approved. We also found that for approved users, the most important things a bank should look at in checking if a customer will be loyal is the number of total transactions and their revolving balance with the bank.

7. References

https://huggingface.co/datasets/Rianknow/creditscoring
https://www.kaggle.com/datasets/sakshigoyal7/credit-card-customers/data
https://www.kaggle.com/datasets/rikdifos/credit-card-approval-prediction
https://www.kaggle.com/datasets/samuelcortinhas/credit-card-approval-clean-data