

NBER WORKING PAPER SERIES

PREDICTING CONSUMER DEFAULT:  
A DEEP LEARNING APPROACH

Stefania Albanesi  
Domonkos F. Vamossy

Working Paper 26165  
<http://www.nber.org/papers/w26165>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
August 2019

This research was supported by the National Science Foundation under Grant No. SES 1824321. This research was also supported in part by the University of Pittsburgh Center for Research Computing through the resources provided. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Stefania Albanesi and Domonkos F. Vamossy. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Predicting Consumer Default: A Deep Learning Approach  
Stefania Albanesi and Domonkos F. Vamossy  
NBER Working Paper No. 26165  
August 2019  
JEL No. C45,C55,D14,D18,E44,G0,G2

**ABSTRACT**

We develop a model to predict consumer default based on deep learning. We show that the model consistently outperforms standard credit scoring models, even though it uses the same data. Our model is interpretable and is able to provide a score to a larger class of borrowers relative to standard credit scoring models while accurately tracking variations in systemic risk. We argue that these properties can provide valuable insights for the design of policies targeted at reducing consumer default and alleviating its burden on borrowers and lenders, as well as macroprudential regulation.

Stefania Albanesi  
Department of Economics  
University of Pittsburgh  
4901 Wesley W. Posvar Hall  
Pittsburgh, PA 15260  
and CEPR  
and also NBER  
stefania.albanesi@gmail.com

Domonkos F. Vamossy  
Department of Economics  
University of Pittsburgh  
230 S Bouquet St  
Pittsburgh, PA 15260  
d.vamossy@pitt.edu

# 1 Introduction

The dramatic growth in household borrowing since the early 1980s has increased the macroeconomic impact of consumer default. Figure 1 displays total consumer credit balances in millions of 2018 USD and the delinquency rate on consumer loans starting in 1985. The delinquency rate mostly fluctuates between 3 and 4%, except at the height of the Great Recession when it reached a peak of over 5%, and in its aftermath when it dropped to a low of 2%. With the rise in consumer debt, variations in the delinquency rate have an ever larger impact on household and financial firm balances sheets. Understanding the determinants of consumer default and predicting its variation over time and across types of consumers can not only improve the allocation of credit, but also lead to important insights for the design of policies aimed at preventing consumer default or alleviating its effects on borrowers and lenders. They are also critical for macroprudential policies, as they can assist with the assessment of the impact of consumer credit on the fragility of the financial system.

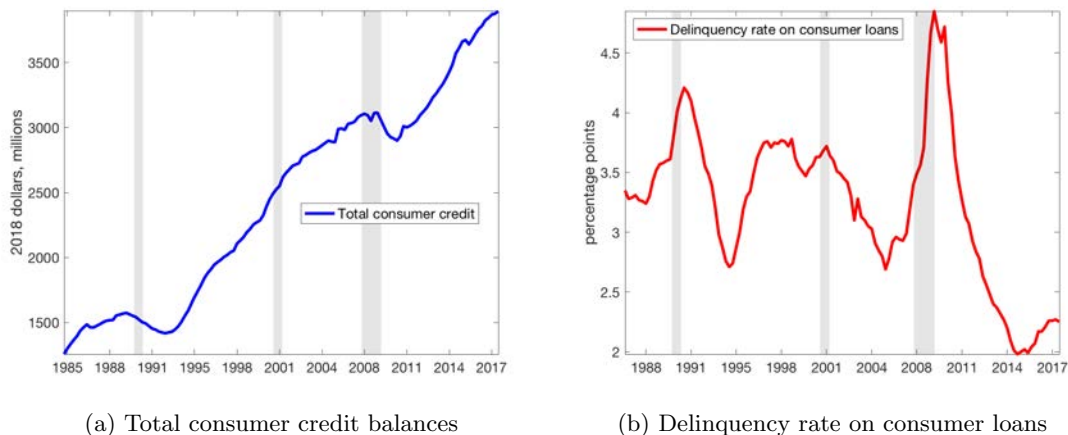


Figure 1: Source: Author’s calculations based on Federal Reserve Board data.

This paper proposes a novel approach to predicting consumer default based on deep learning. We rely on deep learning as this methodology is specifically designed for prediction in environments with high dimensional data and complicated non-linear patterns of interaction among factors affecting the outcome of interest, for which standard regression approaches perform poorly. Our methodology uses the same information as standard credit scoring models, which are one of the most important factors in the allocation of consumer credit. We show that our model substantially improves the accuracy of default predictions while increasing transparency and accountability. It is also able to track variations in systemic risk, and is able to identify the most important factors driving defaults and how they change over time. Finally, we show that adopting our model can accrue substantial savings to borrowers and lenders.

Credit scores constitute one of the most important factors in the allocation of consumer credit

in the United States. They are proprietary measures designed to rank borrowers based on their probability of future default. Specifically, they target the probability of a 90days past due delinquency in the next 24 months.<sup>1</sup> Despite their ubiquitous use in the financial industry, there is very little information on credit scores, and emerging evidence suggests that as currently formulated credit scores have severe limitations. For example, Albanesi, De Giorgi, and Nosal (2017) show that during the 2007-2009 housing crisis there was a marked rise in mortgage delinquencies and foreclosures among high credit score borrowers, suggesting that credit scoring models at the time did not accurately reflect the probability of default for these borrowers. Additionally, it is well known that credit scores are indiscriminately low for young borrowers, and a substantial fraction of borrowers are unscored, which prevents them from accessing conventional forms of consumer credit.

The Fair Credit Reporting Act, a legislation passed in 1970, and the Equal Opportunity in Credit Access Act of 1984 regulate credit scores and in particular determine which information can be included and must be excluded in credit scoring models. Such models can incorporate information in a borrower’s credit report, except age and location. These restrictions are intended to prevent discrimination by age and factors related to location, such as race.<sup>2</sup> The law also mandates that entities that provide credit scores make public the four most important factors affecting scores. In marketing information, these are reported to be length of credit history, which is stated to explain about 15% of variation in credit scores, and credit utilization, number and variety of the debt products and inquiries, each stated to explain about 25-30% of the variation in credit scores. Other than this, there is very little public information on credit scoring models, though several services are now available that allow consumers to simulate how various scenarios, such as paying off balances or taking out new loans, will affect their scores.

The purpose of our analysis is to propose a model to predict consumer default that uses the same data as conventional credit scoring models, improves on their performance, benefiting both lenders and borrowers, and provides more transparency and accountability. To do so, we resort to deep learning, a type of machine learning ideally suited to high dimensional data, such as that available in consumer credit reports.<sup>3</sup> Our model uses inputs as features, such as debt balances and number of trades, delinquency information, and attributes related to the length of a borrower’s credit history, to produce an individualized estimate that can be interpreted as a probability of default. We target the same default outcome as conventional credit scoring models, namely a 90+ days delinquency in the subsequent 8 quarters. For most of the analysis, we train the model on data

---

<sup>1</sup>The most commonly known is the FICO score, developed by the FICO corporation and launched in 1989. The three credit reporting companies or CRCs, Equifax, Experian and TransUnion have also partnered to produce VantageScore, an alternative score, which was launched in 2006. Credit scoring models are updated regularly. More information on credit scores is reported in Section 6.1 and Appendix E.

<sup>2</sup>Credit scoring models are also restricted by law from using information on race, color, gender, religion, marital status, salary, occupation, title, employer, employment history, nationality.

<sup>3</sup>For excellent reviews of how machine learning can be applied in economics, see Mullainathan and Spiess (2017) and Athey and Imbens (2019).

for one quarter and test it on data 8 quarters ahead, in keeping with the default outcome we are considering, so that our predictions are truly out of sample. We present a variety of performance metrics suggesting that our model has very strong predictive ability. Accuracy, that is percent of observations correctly classified, is above 86% for all periods in our sample, and the AUC-Score, a commonly used metric in machine learning, is always above 92%.

To better assess the validity of our approach, we compare our deep learning model to logistic regression and a number of other machine learning models. Deep learning models feature multiple hidden layers, designed to capture multi-dimensional feature interactions. By contrast, logistic regression can be interpreted as a neural network without any hidden layers. Our results suggest that deep learning is necessary to capture the complexity associated with default behavior, since all deep models perform substantially better than logistic regression. The importance of feature interaction reflects the complexity associated with default behavior. Additionally, our optimized model combines a deep neural network and gradient boosting and outperforms other machine learning models, such as random forests and decision trees, as well as deep neural networks and gradient boosting in isolation. However, all approaches show much stronger performance than logistic regression, suggesting that the main advantage is the adoption of a deep framework.

We also compare the performance of our model to a conventional credit score. By construction, credit scores only provide an ordinal ranking of consumers based on their default risk, and are not associated to a specific default probability. Yet, it is still possible to compare performance by assessing whether borrowers fall in different points of the distribution with the credit score compared to our model predictions. We find that our model performs significantly better than conventional credit scores. The rank correlation between realized default rates and the credit score is about 98%, where it is close to 1 for our model. Additionally, the Gini coefficient for the credit score, a measure of the ability to differentiate borrowers based on their credit score is approximately 81% and drops during the 2007-2009 crisis, while the Gini coefficient for our model is approximately 86% and stable over time. Perhaps most importantly, the credit score generates large disparities between the implied predicted probability of default and the realized default rate for large groups of customers, particularly at the low end of the credit score distribution. As an illustration, among Subprime borrowers, 17% display default behavior which is consistent with Near Prime borrowers and 15% display default behavior consistent with Deep Subprime. The default rates for Deep Subprime, Subprime and Near Prime borrowers are respectively 95%, 79% and 44%, so this misclassification is large, and it would imply large losses for lenders and borrowers in terms of missed revenues or higher interest rates. By contrast, the discrepancy between predicted and realized default rates for our model is never more than 4 percentage points for categories with at least a percent share of default risk.

Another advantage of our approach when compared to conventional credit scoring models is that we can generate a predicted probability of default for a much larger class of borrowers. Borrowers may be unscored because they do not have sufficient information in their credit report or because

the information is stale, and approximately 8% of borrowers fall into this category.<sup>4</sup> The absence of a credit score implies that these borrowers do not qualify for most types of credit and is very consequential. Our model can generate a predicted probability of default for all borrowers with a non-empty credit record. We achieve this in part by not including lags in our specification, which implies that only current information in a borrower’s credit report is used. This is not costly from a performance standpoint as many attributes used as inputs in the model are temporal in nature and capture lagged behavior, such as ”worst status on all trades in the last 6 months.”

We also examine the ability of our model to capture the evolution of aggregate default risk. Since our data set is nationally representative and we can score all borrowers with a non-empty credit record, the average predicted probability of default in the population based on our model corresponds to an estimate of aggregate default risk. We find that our model tracks the behavior of aggregate default rates remarkably well. It is able to capture the sharp rise in aggregate default rates in the run up and during the 2007-2009 crisis and also captures the inversion point and the subsequent drastic reduction in this variable. With the growth in consumer credit, household balance sheets have become very important for macroeconomic performance. Having an accurate assessment of the financial fragility of the household sector, as captured by the predicted probability of default on consumer credit has become crucially important and can aid in macro prudential regulation, as well as for designing fiscal and monetary policy responses to adverse aggregate economic shocks. This is another advantage of our model compared to credit scores, since the latter only provides an ordinal ranking of consumers with respect to their probability of default. Our model can provide such a ranking but in addition also provides an individual prediction of the default rate which can be aggregated into a systemic measure of default risk for the household sector.

As a final application, we compute the value to borrowers and lenders of using our model. For consumers, the comparison is made relative to the credit score. Specifically, we compute the credit card interest rate savings of being classified according to our model relative to the credit score. Being placed in a higher default risk category substantially increases the interest rates charged on credit cards at origination and increasingly so as more time lapses since origination, whereas being placed in a lower risk category reduces interest rate costs. We choose credit cards as they are a very popular form of unsecured debt, with 73% of consumers holding at least one credit or bank card. In percentage of credit cards balances, average net interest rate expense savings are approximately 5% for low credit score borrowers. These values constitute lower bounds as they do not include the higher fees and more stringent restrictions associated with credit cards targeted to low credit score borrowers and the increased borrowing limits available to higher credit score borrowers. For lenders, we calculated the value added by using our model in comparison to not having a prediction of default risk or having a prediction based on logistic regression. We use logistic regression for this exercise as it is understood to be the main methodology for conventional credit scoring models. Over a loan with a three year amortization period, we find that the gains relative to no forecast

---

<sup>4</sup>See Brevoort, Grimm, and Kambara (2016). For more information, see Section 6.1.1.

are in the order of 75% with a 15% interest rate, while the gains for relative to a model based on logistic regression are approximately 5%. These results suggest that both borrowers and lenders would experience substantial gains from switching to our model.

Our analysis contributes to the literature on consumer default in a variety of ways. We are the first to develop a prediction model of consumer default using credit bureau data that complies with all of the restrictions mandated by U.S. legislation in this area, and we do so using a large and temporally extended panel of data. This enables us to evaluate model performance in a setting that is closer to the one prevailing in the industry and to train and test our model in a variety of different macroeconomic conditions. Previous contributions either focus on particular types of default or use transaction data that is not admissible in conventional credit scoring models. The closest contributions to our work are Khandani, Kim, and Lo (2010), Butaru et al. (2016) and Sirignano, Sadhwani, and Giesecke (2018). Khandani, Kim, and Lo (2010) apply a decision tree approach to forecast credit card delinquencies with data for 2005-2009. They estimate cost savings of cutting credit lines based on their forecasts and calculate implied time series patterns of estimated delinquency rates. Butaru et al. (2016) apply machine learning techniques to combined consumer trade line, credit bureau, and macroeconomic variables for 2009-2013 to predict delinquency. They find substantial heterogeneity in risk factors, sensitivities, and predictability of delinquency across lenders, implying that no single model applies to all institutions in their data. Sirignano, Sadhwani, and Giesecke (2018) examine over 120 million mortgages between 1995 to 2014 to develop prediction models of multiple states, such as probabilities of prepayment, foreclosure and various types of delinquency. They use loan level and zip code level aggregate information. They also provide a review of the literature using machine learning and deep learning in financial economics. Kvamme et al. (2018) also predict mortgage default using convolutional neural networks and emphasize the advantages of deep learning. Both Sirignano, Sadhwani, and Giesecke (2018) and Kvamme et al. (2018) do not evaluate their models out of sample. Finally, Lessmann et al. (2015) reviews the recent literature on credit scoring, which is based on substantially smaller datasets than the one we have access to, and recommends random forests as a possible benchmark. However, we find that our hybrid model as well as our model components, a deep neural network and gradient boosted trees, improves substantially over random forests, possibly owing to recent methodological advances in deep learning, including the use of dropout, the introduction of new activation functions and the ability to train larger models.

Our model is interpretable, which implies that we are able to assess the most important factors associated with default behavior and how they vary over time. This information is important for lenders, and can be used to comply with legislation that requires lenders and credit score providers to notify borrowers of the most important factors affecting their credit score. Additionally, it can be used to formulate economic models of consumer default. The literature on consumer default<sup>5</sup>

---

<sup>5</sup> Some notable contributions include Chatterjee et al. (2007), Livshits, MacGee, and Tertilt (2007), and Athreya, Tam, and Young (2012).

suggests that the determinants of default are related to preferences, such as impatience which increases the propensity to borrow, or adverse expenditure of income shocks. Based on these theories, it is then possible to construct theoretical models of credit scoring, of which Chatterjee et al. (2016) is a leading example. We find that the number of trades and the balance on outstanding loans are the most important factors associated with an increase in the probability of default, in addition to outstanding delinquencies and length of the credit history. This information can be used to improve models of consumer default risk and enhance their ability to be used for policy analysis and design.

We also identify and quantify a variety of limitations of conventional credit scoring models, particularly their tendency to misclassify borrowers by default risk, especially for relatively risky borrowers. This implies that our default predictions could help improve the allocation of credit in a way that benefits both lenders, in the form of lower losses, and borrowers, in the form of lower interest rates. Our results also speak to the perils associated with using conventional credit scores outside on the consumer credit sphere. As it is well known, credit scores are used to screen job applicants, in insurance applications, and a variety of additional settings. Economic theory would suggest that this is helpful, as long as credit score provide information which is correlated with characteristics that are of interest for the party using the score (Corbae and Glover (2018)). However, as we show, conventional credit scores misclassify borrowers by a very large degree based on their default risk, which implies that they may not be accurate and may not include appropriate information or use adequate methodologies. The broadening use of credit scores would amplify the impact of these limitations.

The paper is structured as follows. Section 2 describes our data. Section 3 discusses the patterns of consumer default that motivate our adoption of deep learning. Section 4 describes our prediction problem and our model. Section 5 provides a comprehensive performance assessment of our model, compares it to other approaches, and uses a variety of interpretability techniques to understand which factors are strongly associated with default behavior. Section 6 compares our model to conventional credit scores, illustrates its performance in predicting and quantifying aggregate default risk and calculates the value added of adopting our model over alternatives for lenders and borrowers.

## 2 Data

We use anonymized credit file data from the Experian credit bureau. The data is quarterly, it starts in 2004Q1 and ends in 2015Q4. The data comprises over 200 variables for an anonymized panel of 1 million households. The panel is nationally representative, constructed from a random draw for the universe of borrowers with an Experian credit report. The attributes available comprise information on credit cards, bank cards, other revolving credit, auto loans, installment loans, business loans, first and second mortgages, home equity lines of credit, student loans and collections. There is



information on the number of trades for each type of loan, the outstanding balance and available credit, the monthly payment, and whether any of the accounts are delinquent, specifically 30, 60, 90, 180 days past due, derogatory or charged off. All balances are adjusted for joint accounts to avoid double counting. Additionally, we have the number of hard inquiries by type of product, and public record items, such as bankruptcy by chapter, foreclosure and liens and court judgments. For each quarter in the sample, we also have each borrowers’s credit score. The data also includes an estimate of individual and household labor income based on IRS data. Because this is data drawn from credit reports, we do not know gender, marital status or any other demographic characteristic, though we do know a borrower’s address at the zip code level. We also do not have any information on asset holdings.

Table 1 reports basic demographic information on our sample, including age, household income, credit score and incidence of default, which here is defined as the fraction of households who report a 90 or more days past due delinquency on any trade. This will be our baseline definition of default, as this is the outcome targeted by credit scoring models. Approximately 34% of consumers display such a delinquency.

Table 1: Descriptive Statistics

Feature	Mean	Std. Dev.	Min	25%	50%	75%	Max
Age	45.8	16.3	18	32.2	45.1	57.8	83
Household Income	77.1	55.0	15	42	64	90	325
Credit Score	678.4	111.0	300	588	692	780	839
Default within 8Q	0.339	0.473	0	0	0	1	1

Credit score corresponds to Vantage Score 3. Household income is in USD thousands, trimmed at the 99th percentile. Source: Authors’ calculations based on Experian Data.

### 3 Patterns in Consumer Default

We now illustrate the complexity of the relation between the various factors that are considered important drivers of consumer default. Our point of departure are standard credit scoring models. While these models are proprietary, the Fair Credit Reporting Act of 1970 and the Equal Opportunity in Credit Access Act of 1984 mandate that the 4 most important factors determining the credit scores be disclosed, together with their importance in determining variation in credit scores. These include credit utilization and number of hard inquiries, which are supposed to capture a consumer’s demand for credit, the variety of debt products, which capture the consumer’s experience in managing credit, and the number and severity of delinquencies. Each of these factors is stated to account for 25-30% of the variation in credit scores. The length of the credit history is also seen

as a proxy on a consumer’s experience in managing credit, and this is reported as accounting for 10-15% of the variation in credit scores.<sup>6</sup> The models used to determine credit scores as a function of these attributes are not disclosed, but they are widely believed to be based on linear and logistic regression as well as score cards. Additionally, available credit scoring algorithms typically do not score all borrowers.

Subsequently, we illustrate the properties of consumer default that suggest deep learning might be a good candidate for developing a prediction model. Specifically, we show that default is a relatively rare but very persistent outcome, there are substantial non-linearities in the relation between default and plausible covariates, as well as high order interactions between covariates and default outcomes.

### 3.1 Default Transitions

The default outcome we consider is a 90+ days delinquency, which occurs if the borrower has missed scheduled payments on any product for 90 days or more.<sup>7</sup> This is the default outcome targeted by the most widely used credit scoring models, which rank consumers based on their probability of becoming 90+ days delinquent in the subsequent 8 quarters. We refer to borrowers who are either current or up to 60 days delinquent on their payments as *current*.

The transition matrix from current to 90+ days past due in the subsequent 8 quarters is given in Table 2. Clearly, the two states are both highly persistent, with a 77% of current customers remaining current in the next 8 quarters, and 93% of customers in default remaining in that state over the same time period. The probability of transition from current to default is 23%, while the probability of curing a delinquency with a transition from default to current is only 7%. These results suggest that default is a particularly persistent state, and predicting a transition into default is very valuable from the lender’s perspective, since they are unlikely to be able to recuperate their losses. But it is also quite difficult, as the current state is also very persistent.

### 3.2 Non-linearities

Our model includes a relatively large list of features, which is presented in Table 17. The summary statistics for these features are reported in Table 18 in the Appendix. As is demonstrated in the table, there is a wide dispersion in the distribution of these variables. For example, the average balance on credit card trades is approximately \$4,500, but the standard deviation, at \$9,800, is more than twice as large. Similarly, average total debt balances are approximately \$77,000, while the standard deviation is \$170,000 and the 75th percentile \$95,000, suggesting a high upper tail

---

<sup>6</sup>For an overview of the information available to borrowers about the determinants for their credit score, see <https://ficoscore.com>.

<sup>7</sup>For instance, if no payment has been made by the last day of the month within the past three months and the payment was due on the first day of the month three months ago. For credit cards, this occurs if the borrower does not make at least their minimum payment.

Table 2: Default Transitions

Current/Next 8Q	No default	Default
No default	0.776	0.224
Default	0.073	0.927

Quarterly frequency of transition from current to default. Current corresponds to 0-89 day past due on any trade, Default corresponds to 90+ day past due on any trade in the subsequent 8 quarters. Source: Authors' calculations based on Experian Data.

dispersion of this variable. The other features display similar patterns.

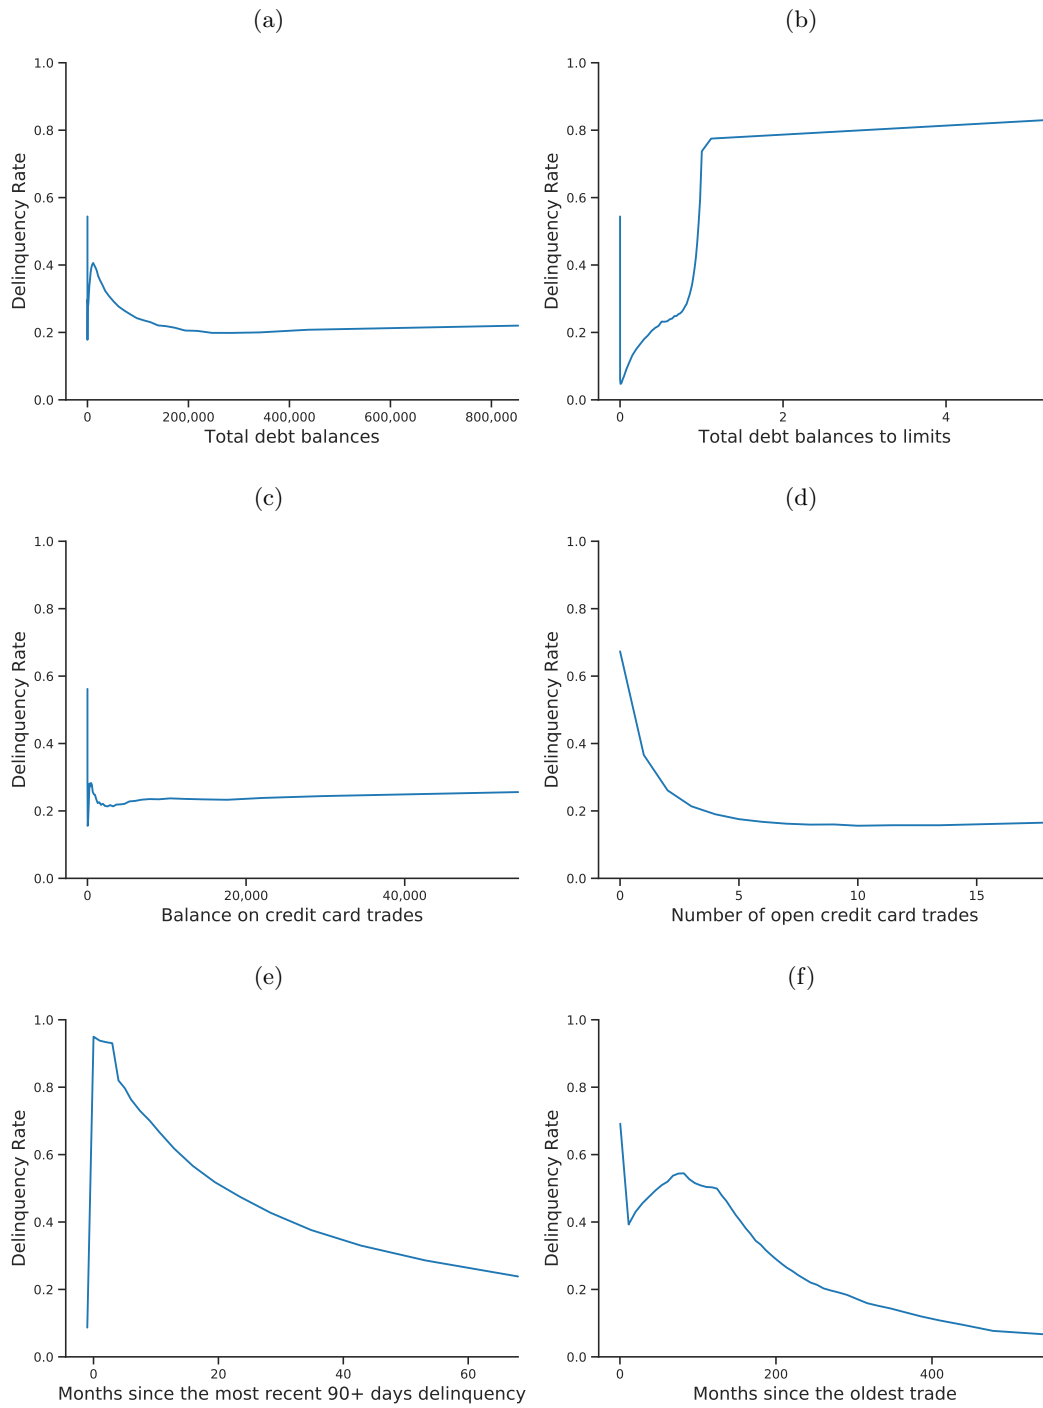
The features are used to predict the probability of default. We now illustrate the highly non-linear relation between the features and the incidence of default. Figure 2 shows how the default rate, defined as the fraction of borrowers with a 90+ day past due delinquency in the subsequent 8 quarters, varies with total debt balances, credit utilization, the credit limit on credit cards, the number of open credit card trades, the number of months since the most recent 90+ day past due delinquency and the months since the oldest trade was opened. The figures show that while the relation between the features and the incidence of default is mostly monotone, it is highly nonlinear, with very little variation in the incidence of default for most intermediate values of the variable and much higher or lower values at the extremes of the range of each covariate. The variables in the figure are just illustrative, a similar pattern holds for most plausible features.

### 3.3 High Order Interactions

Multidimensional interactions are another feature of the relation between default and plausible covariates, that is default behavior is simultaneously related with multiple variables. To see this, Figure 3 presents contour plots of the relation between the incidence of default and couples of covariates. The covariates reported here are chosen since they are important driving factors in default decisions, based on our model, as discussed in Section 5.2. Panels (a) and (b) explore the joint variation in the incidence of default with total debt balances, credit utilization (total debt balances to limits), and credit history. Blue values correspond to high delinquency rates while red values to low delinquency rates. As can be seen from both panels, higher credit utilization corresponds to higher delinquency rate, but for given credit utilization, an increase in total debt balances first decreases then increases the delinquency rate, where the switch in sign depends on the utilization rate. For given utilization rates, a longer credit history first increases then decreases the delinquency rate, provided the utilization rate is smaller than 1.<sup>8</sup> Panels (c) and (d) explore the relation between default and credit card borrowing. Default rates decline with the number of

<sup>8</sup>Utilization rates above 1 can arise for a delinquent borrower if fees and other penalty add to their balances for given credit limits.

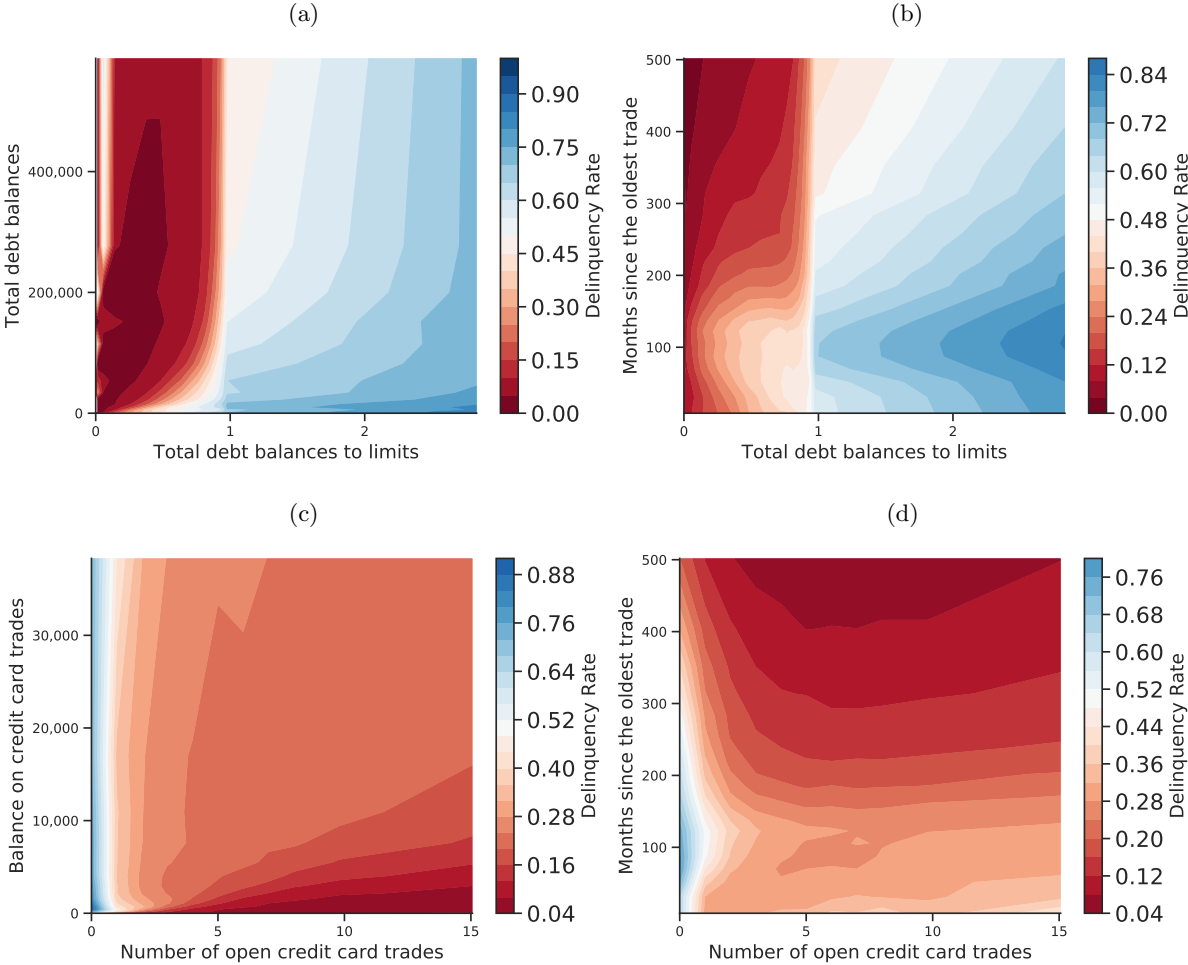
Figure 2: Nonlinear Relation Between Default and Covariates



Delinquency rate is the fraction with 90+ days past due trades in subsequent 8 quarters. In panel (e) and (f), -1 implies no past delinquency. Source: Authors' calculations based on Experian Data.

credit cards, though for a given number of credit card trades, they mostly increase with credit card balances. This relation, however varies with the level of both variables. An increase in the length of credit history is typically associated with lower default rates, however, if the number of open credit cards is low, this relation is non-monotone. The variables reported in the figures are illustrative of a general pattern in the joint relation between couples of covariates and default rates.

Figure 3: Multidimensional Relation Between Default and Covariates



Relationship between 90+ days past due delinquency rate and pairs of covariates. Source: Authors' calculations based on Experian Data.

This pattern of multidimensional non-linear interactions across covariates is fairly difficult to model using standard econometric approaches. For this reason, we propose a deep learning approach to be explained below.

## 4 Model

Predicting consumer default maps well into a supervised learning framework, which is one of the most widely used techniques in the machine learning literature. In supervised learning, a learner takes in pairs of input/output data. The input data, which is typically a vector, represent pre-identified attributes, also known as features, that are used to determine the output value. Depending on the learning algorithm, the input data can contain continuous and/or discrete values with or without missing data. The supervised learning problem is referred to as a "regression problem" when the output is continuous, and as a "classification problem" when the output is discrete. Once the learner is presented with input/output data, its task is to find a function that maps the input vectors to the output values. A brute force way of solving this task is to memorize all previous values of input/output pairs. Though this perfectly maps the input data to the output values in the training data set, it is unlikely to succeed in forecasting the output values if (1) the input values are different from the ones in the training data set or (2) when the training data set contains noise. Consequently, the goal of supervised learning is to find a function that generalizes beyond the training set, so that it correctly forecasts out-of-sample outcomes. Adopting this machine-learning methodology, we build a model that predicts defaults for individual consumers. We define default as a 90+ days delinquency on any debt in the subsequent 8 quarters, which is the outcome targeted by conventional credit scoring models. Our model outputs a continuous variable between 0 and 1 that can be interpreted under certain conditions as an estimate of the probability of default for a particular borrower at a given point in time, given input variables from their credit reports.

We start by formalizing our prediction problem. We adopt a discrete-time formulation for periods  $0, 1, \dots, T$ , each corresponding to a quarter. We let the variable  $D_t^i$  prescribe the state at time  $t$  for individual  $i$  with  $D \subset \mathbb{N}$  denoting the set of states. We define  $D_1^i = 1$  if a consumer is 90+ days past due on any trade and  $D_1^i = 0$  otherwise. Consumers will transition between these two states over their lifetime.

Our target outcome is 90+ days past due in the subsequent 8 quarters, defined as:

$$Y_t^i = \begin{cases} 0 & \text{if } \sum_{n=t}^{t+7} D_n^i = 0 \\ 1 & \text{otherwise} \end{cases} \quad (1)$$

We allow the dynamics of the state process to be influenced by a vector of explanatory variables  $X_{t-1}^i \in \mathbb{R}^{d_x}$ , which includes the state  $D_{t-1}^i$ . In our empirical implementation,  $X_{t-1}^i$  represents the features in Table 17. We fix a probability space  $(\Omega, \mathcal{F}, \mathbb{P})$  and an information filtration  $(\mathcal{F}_t)_{(t=0,1,\dots,T)}$ . Then, we specify a probability transition function  $h_\theta : \mathbb{R}^{d_x} \rightarrow [0, 1]$  satisfying

$$\mathbb{P}[Y_t^i = y | \mathcal{F}_{t-1}] = h_\theta(X_{t-1}^i), y \in D \quad (2)$$

where  $\theta$  is a parameter to be estimated. Equation 2 gives the marginal conditional probability for

the transition of individual  $i$ 's debt from its state  $D_{t-1}^i$  at time  $t - 1$  to state  $y$  at time  $t$  given the explanatory variables  $X_{t-1}^i$ .<sup>9</sup> Let  $g$  denote the standard *softmax* function:

$$g(z) = \left( \frac{1}{1 + e^{-z}} \right), z \in \mathbb{R}^K, \quad (3)$$

where  $K = |D|$ . The vector output of the function  $g$  is a probability distribution on  $D$ .

The marginal probability defined in equation 2 is the theoretical counterpart of the empirical transition matrix reported in Table 2. We propose to model the transition function  $h_\theta$  with a hybrid deep neural network/gradient boosting model, which combines the predictions of a deep neural network and an extreme gradient boosting model. We explain each of the component models and their properties and the rationale for combining them below.

#### 4.1 Deep Neural Network

One component of our model is based on deep learning, in the class used by Sirignano, Sadhwani, and Giesecke (2018). We restrict attention to feed-forward neural networks, composed of an input layer, which corresponds to the data, one or more interacting hidden layers that non-linearly transform the data, and an output layer that aggregates the hidden layers into a prediction. Layers of the networks consist of neurons with each layer connected by synapses that transmit signals among neurons of subsequent layers. A neural network is in essence a sequence of nonlinear relationships. Each layer in the network takes the output from the previous layer and applies a linear transformation followed by an element-wise non-linear transformation.

Figure 4 illustrates an example of a two layer neural network. This neural network has 3 input units (denoted  $x_1, x_2, x_3$ ), 4 hidden units, and 1 output unit. Let  $n_l$  denote the number of layers in this network ( $n_l = 2$ ). We label layer  $l$  as  $L_l$ , where layer  $L_0$  is the input layer, and layer  $L_{L=2}$  is the output layer. The layers between the input ( $l = 0$ ) and the output layer ( $l = L$ ) are called hidden layers. Given this notation, there are  $L - 1$  hidden layers, 1 in this specific example. A neural network without any hidden layers ( $L = 1$ ) is a logistic regression model.

There are two ways to increase the complexity a neural network: (1) increase the number of hidden layers and (2) increase the number of units in a given layer. Lower tier layers in the neural network learn simpler patterns, from which higher tier layers learn to produce more complex patterns. Given a sufficient number of neurons, neural networks can approximate continuous functions on compact sets arbitrarily well (see Hornik, Stinchcombe, and White (1989) and Hornik (1991)). This includes approximating interactions (i.e., the product and division of features). There are two main advantages of adding more layers over increasing the number of units to existing layers; (1) later layers build on early layers to learn features of greater complexity and (2) deep neural networks— those with three or more hidden layers— need exponentially fewer neurons than shallow

---

<sup>9</sup>The state  $y$  encompasses realizations of the state between time  $t$  and  $t + 7$ .




Figure 4: Two Layer Neural Network Example

networks (Bengio and LeCun (2007) and Montúfar et al. (2014)).

In the neural network represented in Figure 4, the parameters to be estimated are  $(W, b) = (W^{(0)}, b^{(0)}, W^{(1)}, b^{(1)})$ , where  $W_{ij}^{(l)}$  denotes the weight associated with the connection between unit  $j$  in layer  $l$  and unit  $i$  in layer  $l + 1$ , and  $b_i^{(l)}$  is the bias associated with unit  $i$  in layer  $l + 1$ . Thus, in this example  $W^{(0)} \in \mathbb{R}^{3 \times 4}$ ,  $b^{(0)} \in \mathbb{R}^{4 \times 1}$  and  $W^{(1)} \in \mathbb{R}^{1 \times 4}$ ,  $b^{(1)} \in \mathbb{R}$ . This implies that there are a total of  $21 = (3+1)*4+5$  parameters (four parameters to reach each neuron and five weights to aggregate the neurons into a single output). In general, the number of weight parameters in each hidden layer  $l$  is  $N^{(l)}(1 + N^{(l-1)})$ , plus  $1 + N^{(L-1)}$  for the output layer, where  $N^{(l)}$  denotes the number of neurons in each layer  $l = 1, \dots, L$ .

Let  $a_i^{(l)}$  denote the activation (e.g., output value) of unit  $i$  in layer  $l$ . Fix  $W$  and  $b$ , our neural network defines a hypothesis  $h_{W,b}(x)$  that outputs a real number between 0 and 1.<sup>10</sup> Let  $f(\cdot)$  denote the activation function that applies to vectors in an element-wise fashion. The computation this neural network represents, often referred to as forward propagation, can be written as:

$$z^{(1)} = W^{(0),T}x + b^{(0)}$$

$$a^{(1)} = f(z^{(1)})$$

$$z^{(2)} = W^{(1),T}a^{(1)} + b^{(1)}$$

$$h_{W,b}(x) = a^{(2)} = f(z^{(2)})$$

<sup>10</sup>This is a property of the sigmoid activation function.



There are many choices to make when structuring a neural network, including the number of hidden layers, the number of neurons in each layer, and the activation functions. We built a number of network architectures having up to fifteen hidden layers.<sup>11</sup> All architectures are fully connected so each unit receives an input from all units in the previous layer.

Neural networks tend to be low-bias, high-variance models, which imparts them a tendency to over-fit the data. We apply dropout to each of the layers to avoid over-fitting (see Srivastava et al. (2014)). During training, neurons are randomly dropped (along with their connections) from the neural network with probability  $p$  (referred to as the dropout rate), which prevents complex co-adaptations on training data.

We apply the same activation function (scaled exponential linear unit or SELU) at all nodes, which is obtained via hyperparameter optimization,<sup>12</sup> and defined as:

$$\text{SELU}(x) = \lambda \begin{cases} \alpha e^x - \alpha & \text{if } x \leq 0 \\ x & \text{otherwise} \end{cases} \quad (4)$$

where  $\alpha \approx 1.6733$ , and  $\lambda \approx 1.0507$ .

Let  $N^{(l)}$  denote the number of neurons in each layer  $l = 1, \dots, L$ . Define the output of neuron  $k$  in layer  $l$  as  $z_k^{(l)}$ . Then, define the vector of outputs (including the bias term  $z_0^{(l)}$ ) for this layer as  $z^{(l)} = (z_0^{(l)}, z_1^{(l)}, \dots, z_{N^{(l)}}^{(l)})'$ . For the input layer, define  $z^{(0)} = (x_0^{(l)}, x_1^{(l)}, \dots, x_{N^{(l)}}^{(l)})'$ . Formally, the recursive output of the  $l$ -th layer of the neural network is:

$$z^{(l)} = \text{SELU}(W^{(l-1),T} z^{(l-1)} + b^{(l-1)}), \quad (5)$$

with final output:

$$h_\theta(x) = g(W^{(L-1),T} z^{(L-1)} + b^{(L-1)}). \quad (6)$$

The parameter specifying the neural network is:

$$\theta = (W_0, b_0, \dots, W_{L-1}, b_{L-1}) \quad (7)$$

## 4.2 Decision Tree Models

The second component of our model is Extreme Gradient Boosting, which builds on decision tree models. Tree-based models split the data several times based on certain cutoff values in the explanatory variables.<sup>13</sup> A number of such models have become quite prevalent in the literature, most notably random forests (see Breiman (2001) and Butaru et al. (2016)) and Classification and Regression Trees, known as CART. We briefly review CART and then explain gradient boosting.

<sup>11</sup>The number of layers and the number of neurons in each layer, along with other hyperparameters of the model, are chosen by Tree-structured Parzen Estimator (TPE) approach. See Appendix C for more details.

<sup>12</sup>There are many potential choices for the nonlinear activation function, including the sigmoid, relu, and tanh.

<sup>13</sup>Splitting means that different subsets of the dataset are created, where each observation belongs to one subset. For a review on decision trees, see Khandani, Kim, and Lo (2010).

### 4.2.1 CART

There are a number of different decision tree-based algorithms. As an illustration of the approach, we describe Classification and Regression Trees or CART. CART models an outcome  $y_i$  for an instance  $i$  as follows:

$$\hat{y}_i = \hat{f}(x_i) = \sum_{m=1}^M c_m I\{x_i \in R_m\}, \quad (8)$$

where each observation  $x_i$  belongs to exactly one subset  $R_m$ . The identity function  $I$  returns 1 if  $x_i$  is in  $R_m$  and 0 otherwise. If  $x_i$  falls into  $R_l$ , the predicted outcome is  $\hat{y} = c_l$ , where  $c_l$  is the mean of all training observations in  $R_l$ .

The estimation procedure takes a feature and computes the cut-off point that minimizes the Gini index of the class distribution of  $\mathbf{y}$ , which makes the two resulting subsets as different as possible. Once this is done for each feature, the algorithm uses the best feature to split the data into two subsets. The algorithm is then repeated until a stopping criterium is reached.

Tree-based models have a number of advantages that make them popular in applications. They are invariant to monotonic feature transformations and can handle categorical and continuous data in the same model. Like deep neural networks, they are well suited to capturing interactions between variables in the data. Specifically, a tree of depth  $L$  can capture  $(L - 1)$  interactions. The interpretation is straightforward, and provides immediate counterfactuals: "If feature  $x_j$  had been bigger / smaller than the split point, the prediction would have been  $\bar{y}_0$  instead of  $\bar{y}_1$ ." However, these models also have a number of limitations. They are poor at handling linear relationships, since tree algorithms rely on splitting the data using step functions, an intrinsically non-linear transformation. Trees also tend to be unstable, so that small changes in the training dataset might generate a different tree. They are also prone to overfitting to the training data. For more information on tree-based models see Molnar (2019).

### 4.2.2 eXtreme Gradient Boosting (XGBoost)

Gradient Boosted Trees (GBT) are an ensemble learning method that corrects for tree-based models' tendency to overfit to training data by recursively combining the forecasts of many over-simplified trees. Though shallow trees are "weak learners" on their own with little predictive power, the theory behind boosting proposes that a collection of weak learners, as an ensemble, creates a single strong learner with improved stability over a single complex tree.

At each step  $m$ ,  $1 \leq m \leq M$ , of gradient boosting, an estimator,  $h_m$ , is computed on the residuals from the previous models predictions. A critical part of gradient boosting method is regularization by shrinkage as proposed by Friedman (2001). This consists in modifying the update rule as follows:

$$F_m(x) = F_{m-1}(x) + \nu \gamma_m h_m(x), \quad (9)$$

where  $h_m(x)$  represents a weak learner of fixed depth,  $\gamma_m$  is the step length and  $\nu$  is the learning

rate or shrinkage factor.

XGBoost is a fast implementation of Gradient Boosting, which has the advantages of fast speed and high accuracy. For classification, XGBoost combines the principles of decision trees and logistic regression, so that the output of our XGBoost model is a number between 0 and 1. For the remainder of the paper we refer to XGBoost as GBT.<sup>14</sup>

### 4.3 Hybrid DNN-GBT Model

We examined two techniques to create a hybrid DNN-GBT ensemble model. Ensemble models combine multiple learning algorithms to generate superior predictive performance than could be obtained from any of the constituent learning algorithms alone. The first method combines the two models by replacing the final layer of the neural network with a gradient boosted trees model. Examples of this approach are Chen, Lundberg, and Lee (2018) and Ren et al. (2017). The second, uses both models separately and then averages out the final predicted probabilities of the two models. We found the latter to perform better on our dataset. This method is similar to Kvamme et al. (2018), who combined a convolutional neural network with a random forest by averaging. Thus, our methodology relies on combining the output of the deep neural network with the output of a gradient boosted trees model. This is achieved in two steps:

1. For each observation, run DNN and GBT separately and obtain predicted probabilities for each of the models;
2. Take the arithmetic mean of the predicted probabilities.

## 5 Implementation

Table 17 lists the features from the credit report data we use as inputs in the model. These covariates are chosen based on economic theory (see for example Chatterjee et al. (2007)) as well as based on information from currently used credit scoring models. They include information on balances and credit limits for different types of consumer debt, severity and number of delinquencies, credit utilization by type of product, public record items such as bankruptcy filings by chapter and foreclosure, collection items, and length of the credit history. In order to be consistent with the restrictions of the Fair Credit Reporting Act on 1970 and the Equal Opportunity in Credit Access Act of 1984 we do not include information on age or zip code, and we do not include any information on income, to be consistent with current credit scoring models. Table 17 lists the full set of features used in our machine learning models. It is important to note that we do not use any lagged features. This is because many of the features have a temporal dimension, for example "worst present status

---

<sup>14</sup>For more on XGBoost, see Chen, Lundberg, and Lee (2018) and Ren et al. (2017).

on any traded in the last 6 months.” Importantly, excluding lags enables us to provide a default prediction to any borrower with a non-empty credit record, which implies that we can score virtually all consumers.

## 5.1 Classifier Performance

In this section, we describe the performance of our hybrid model under various training and testing windows. First, we evaluate our model on the pooled sample (2004Q1-2013Q4), where we apply a random 60%-20%-20% split to our training, validation, and testing sets. Then, to account for look-ahead bias, we train and test our models based on 8 quarter windows that were observable at the time of forecast. In particular, we require our training and testing sets to be separated by 8 quarters to avoid overlap. For instance, the second out-of-sample model was calibrated using input data from 2004Q2, from which the parameter estimates were applied to the input data in 2006Q2 to generate forecasts of delinquencies over the 8 quarter window from 2006Q3-2008Q2. This gives us a total of 32+1 calibration and testing periods reported in Table 3. The percentage of 90+ days past due accounts within 8 quarters varies from 32.5% to 35.9%.

The hybrid model outputs a continuous variable that, under certain circumstances, can be interpreted as an estimate of the probability of an account becoming 90+ days delinquent during the subsequent 8 quarters. One measure of the model’s success is its ability to differentiate between accounts that did become delinquent and those that did not; if these two groups have the same forecasts, the model provides no value. Table 3 presents the average forecast for accounts that did and did not fall into the 90+ days delinquency category over the 32+1 evaluation periods. For instance, during the testing period for 2010Q4, the model’s average prediction among the 35.44% of accounts that became 90+ days delinquent was 73.35%, while the average prediction among the 64.56% of accounts that did not was 15.86%. We should highlight that these are truly out-of-sample predictions, since the model is calibrated using input data from 2008Q4. This shows the forecasting power of our model in distinguishing between accounts that will and will not become delinquent within 8 quarters. Furthermore, this forecasting power seems to be stable over the 32+1 calibration and evaluation periods, partly driven by the frequent re-calibration of the model that captures some of the changing dynamics of consumer behavior.

We also look at accounts that are current as of the forecast date but become 90+ days delinquent within the subsequent 8 quarters. In particular, we contrast the model’s average prediction among individuals who were current on their accounts but became 90+ days delinquent with the average prediction among customers who were current and did not become delinquent. Given the difficulty of predicting default among individuals that currently show no sign of delinquency, we anticipate the model’s performance to be less impressive than the values reported in Table 3. Nonetheless, the values reported in Table 4 indicate that the model is able to distinguish between these two populations. For instance, using input data from 2008Q4, the average model prediction for individuals who were current on their debts and became 90+ days delinquent is 45.13%, contrasted with

Table 3: 1 Quarter Ahead Predictions, Full Sample– Hybrid DNN-GBT

Training Window	Testing Window	Data	Predicted	Delinquents	Non-Delinquents
2004Q1-2013Q4	2004Q1-2013Q4	0.3396	0.3363	0.7583	0.1181
2004Q1	2006Q1	0.3248	0.2924	0.6499	0.1204
2004Q2	2006Q2	0.3274	0.3041	0.6759	0.123
2004Q3	2006Q3	0.3306	0.3111	0.6852	0.1263
2004Q4	2006Q4	0.3347	0.3158	0.6861	0.1295
2005Q1	2007Q1	0.341	0.316	0.6854	0.1249
2005Q2	2007Q2	0.3444	0.3202	0.6872	0.1274
2005Q3	2007Q3	0.3469	0.3197	0.6871	0.1246
2005Q4	2007Q4	0.3505	0.3287	0.6958	0.1306
2006Q1	2008Q1	0.3535	0.3373	0.7098	0.1336
2006Q2	2008Q2	0.3545	0.3338	0.7012	0.1321
2006Q3	2008Q3	0.3558	0.336	0.7048	0.1324
2006Q4	2008Q4	0.3587	0.3407	0.7089	0.1348
2007Q1	2009Q1	0.3588	0.3488	0.7213	0.1403
2007Q2	2009Q2	0.358	0.3486	0.7226	0.1401
2007Q3	2009Q3	0.3573	0.3542	0.7296	0.1455
2007Q4	2009Q4	0.3589	0.3552	0.7307	0.145
2008Q1	2010Q1	0.3589	0.3563	0.7293	0.1474
2008Q2	2010Q2	0.3568	0.362	0.7351	0.155
2008Q3	2010Q3	0.3559	0.3644	0.7367	0.1587
2008Q4	2010Q4	0.3544	0.3623	0.7335	0.1586
2009Q1	2011Q1	0.3541	0.3582	0.7286	0.1552
2009Q2	2011Q2	0.3511	0.3589	0.7271	0.1596
2009Q3	2011Q3	0.35	0.3543	0.7217	0.1564
2009Q4	2011Q4	0.3484	0.3544	0.7255	0.1561
2010Q1	2012Q1	0.3467	0.3539	0.7306	0.154
2010Q2	2012Q2	0.3434	0.3504	0.7254	0.1543
2010Q3	2012Q3	0.3396	0.3488	0.7297	0.1529
2010Q4	2012Q4	0.3358	0.3463	0.7286	0.1531
2011Q1	2013Q1	0.3341	0.3475	0.7343	0.1534
2011Q2	2013Q2	0.3317	0.3441	0.7277	0.1537
2011Q3	2013Q3	0.3298	0.3425	0.7328	0.1504
2011Q4	2013Q4	0.3275	0.3399	0.7328	0.1486

Performance metrics for our model of default risk over 32+1 testing windows. For each testing window, the model is calibrated on data over the period specified in the training window, and predictions are based on the data available as of the data in the training window. For example, the fourth row reports the performance of the model calibrated using input data available in 2004Q3, and applied to 2006Q3 data to generate forecasts of delinquencies for within 8 quarter delinquencies. Average model forecasts over all customers, and customers that (ex-post) did and did not become 90+ days delinquent over the testing window are also reported. Source: Authors' calculations based on Experian Data.

12.34% for those who did not. As in Table 3, the model’s ability to distinguish between these two classes is consistent across the 32+1 evaluation periods listed in Table 4.

Under certain conditions, the forecasts generated by our model can be converted to binary decisions by comparing the forecast to a specified threshold and classifying accounts with scores exceeding that threshold as high-risk. Setting the threshold level comes with a trade-off. A low level threshold leads to many accounts being classified as high risk, and even though this approach may accurately capture customers who are actually high-risk and about to default on their payments, it can also give rise to many low-risk accounts incorrectly classified as high-risk. By contrast, a high threshold can result in too many high-risk accounts being classified as low-risk.

This type of trade-off is inherent in any classification problem, and involves trading off Type-I (false positives) and Type-II (false negatives) errors in a classical hypothesis testing context. In the credit risk management context, a cost/benefit analysis can be formulated contrasting false positives to false negatives to make this trade-off explicit, and applying the threshold that will optimize an objective function in which costs and benefits associated with false positives and false negatives are inputs.

A commonly used performance metric in the machine learning and statistics literature is a  $2 \times 2$  contingency table, often referred to as the *confusion matrix*, that describes the statistical behavior of any classification algorithm. In our application, the two rows correspond to ex post realizations of the two types of accounts in our sample, *no default* and *default*. We define *no default* accounts as those who do not become 90+ days delinquent during the forecast period, and *default* accounts as those who do. The two columns correspond to ex ante classifications of the accounts into these categories. If a predictive model is applied to a set of accounts, each account falls into one of the four cells in the confusion matrix, thus the performance of the model can be assessed by the relative frequencies of the entries. In the Neymann-Pearson hypothesis-testing framework, the lower-left entry is defined as Type-I error and the upper right as Type-II error, while the objective of the researcher is to minimize Type-II error (i.e., maximize "power") subject to a fixed level of Type-I error (i.e., "size").

As an illustration, figure 5 shows the confusion matrix for our hybrid DNN-GBT model calibrated using 2011Q4 data and evaluated on 2013Q4 data and a threshold of 50%. This means that accounts with estimated delinquency probabilities greater than 50% are classified as default and 50% or below as no default. For this quarter, the model classified  $61.29\% + 7.19\% = 68.48\%$  of the accounts as no default, of which 61.29% did indeed not default and 7.19% actually defaulted, that is, they were 90+ days delinquent in the subsequent 8 quarters. By the same token, of the  $5.96\% + 25.56\% = 31.52\%$  borrowers who defaulted, the model accurately classified 25.56%. Thus, the model’s accuracy, defined as the percent of instances correctly classified, is the sum of the entries on the diagonal of the confusion matrix, that is,  $61.29\% + 25.56\% = 86.85\%$ .

We can compute three additional performance metrics from the entries of the confusion matrix, which we describe heuristically here and define formally in the appendix. *Precision* measures the

Table 4: 1 Quarter Ahead Predictions, Current– Hybrid DNN-GBT

Training Window	Testing Window	Data	Predicted	Delinquent	Non-delinquent
2004Q1-2013Q4	2004Q1-2013Q4	0.1676	0.1618	0.5421	0.085
2004Q1	2006Q1	0.1844	0.1558	0.4254	0.0948
2004Q2	2006Q2	0.1702	0.1452	0.3967	0.0937
2004Q3	2006Q3	0.1695	0.1487	0.4019	0.097
2004Q4	2006Q4	0.1727	0.151	0.4029	0.0984
2005Q1	2007Q1	0.1805	0.151	0.4002	0.0961
2005Q2	2007Q2	0.1813	0.1527	0.3947	0.0991
2005Q3	2007Q3	0.1831	0.1494	0.3899	0.0955
2005Q4	2007Q4	0.1847	0.1541	0.3995	0.0985
2006Q1	2008Q1	0.189	0.162	0.4172	0.1025
2006Q2	2008Q2	0.1896	0.16	0.4073	0.1022
2006Q3	2008Q3	0.1872	0.1589	0.4055	0.1021
2006Q4	2008Q4	0.1817	0.1574	0.401	0.1034
2007Q1	2009Q1	0.1781	0.1634	0.4182	0.1082
2007Q2	2009Q2	0.1752	0.1626	0.4177	0.1085
2007Q3	2009Q3	0.1713	0.1673	0.4317	0.1127
2007Q4	2009Q4	0.1661	0.1641	0.4232	0.1125
2008Q1	2010Q1	0.1683	0.1689	0.433	0.1154
2008Q2	2010Q2	0.1668	0.1768	0.449	0.1223
2008Q3	2010Q3	0.1661	0.1806	0.4577	0.1254
2008Q4	2010Q4	0.1644	0.1773	0.4513	0.1234
2009Q1	2011Q1	0.1674	0.1789	0.4594	0.1226
2009Q2	2011Q2	0.1668	0.1798	0.4593	0.1239
2009Q3	2011Q3	0.1669	0.1765	0.452	0.1214
2009Q4	2011Q4	0.1597	0.1726	0.4462	0.1206
2010Q1	2012Q1	0.1604	0.1701	0.4458	0.1174
2010Q2	2012Q2	0.1622	0.1702	0.4465	0.1167
2010Q3	2012Q3	0.1598	0.167	0.4435	0.1143
2010Q4	2012Q4	0.1575	0.165	0.4416	0.1133
2011Q1	2013Q1	0.1606	0.1703	0.4591	0.115
2011Q2	2013Q2	0.1603	0.1707	0.4556	0.1163
2011Q3	2013Q3	0.1578	0.1652	0.4531	0.1113
2011Q4	2013Q4	0.1548	0.1615	0.4457	0.1095

Performance metrics for our model of default risk over 32+1 testing windows for customers who are current as of the forecast date but become 90+ days delinquent in the following 8 quarters. For each testing window, the model is calibrated on data over the period specified in the training window columns, and predictions are based on the data available as of the data in the training window. For example, the fourth row reports the performance of the model calibrated using input data available in 2004Q3, and applied to 2006Q3 data to generate forecasts of delinquencies for within 8 quarter delinquencies. Average model forecasts over all current customers, and all current customers that did and did not become 90+ days delinquent over the testing window are also reported. Source: Authors' calculations based on Experian Data.




Figure 5: Confusion matrix for our model of default risk. Rows correspond to actual states, with default defined as 90+ days delinquent, no default otherwise. Classifier threshold: 50%. The numerical example is based on the model calibrated on 2011Q4 data and applied to 2013Q4 to generate out-of-sample predictions. Source: Authors’ calculations based on Experian Data.

model’s accuracy in instances that are classified as default. *Recall* refers to the number of accounts that defaulted as identified by the model divided by the actual number of defaulting accounts. Finally, the *F-measure* is simply the harmonic mean of precision and recall. In an ideal scenario, we would have very high precision and recall.

We can track the trade-off between true and false positives by varying the classification threshold of our model, and this trade-off is plotted in Figure 6. The blue line, called the Receiver Operating Characteristic (ROC) curve, is the pairwise plot of true and false positive rates for different classification thresholds (green line), and as the threshold decreases, the figure shows that the true positive rate increases, but so does the false positive rate. The ROC curve illustrates the non-linear nature of the trade-offs, implying that increase in true positive rates is not always proportionate with the increase in false positive rates. The optimal threshold then considers the cost of false positives with respect to the gain of true positives. If these are equal, the optimal threshold will correspond to the tangent point of the ROC curve with the 45 degree line.

The last performance metric we consider is the area under the ROC curve, known as AUC score, which is a widely used measure in the machine-learning literature for comparing models. It can be interpreted as the probability of the classifier assigning a higher probability of being in default to an account that is actually in default. The ROC area of our model ranges from 0.9241 to 0.9306, demonstrating that our machine-learning classifiers have strong predictive power in separating the two classes.

Table 5 reports the performance metrics widely used in the machine-learning literature for






Figure 6: Receiver Operating Characteristic (ROC) curve and corresponding classification threshold value of out-of-sample forecasts of 90+ days delinquencies over the 8Q forecast horizon based on our model of default risk. The model is trained on 2011Q4 input data and evaluated on data from 2013Q4. Source: Authors’ calculations based on Experian Data.

each of the 32+1 models discussed. Our models exhibit strong predictive power across the various performance metrics. For instance, the 85.31% precision implies that when our classifier predicts that someone is going to default, there is an 85.31% chance this person will actually default; while the 70.16% recall means that we accurately identified 70.16% of all the defaulters. Our approach of using only one quarter of data to train the model is rather restrictive. Using more quarters usually increases model performance, so since most credit scoring applications will use a training data that exceeds one quarter, performance metrics are likely to improve relative to what we report in our exercise.

Table 6 reports the same performance metrics for the population of borrowers who are current, that is, they do not have any delinquencies in the quarter they are assessed. As previously noted, this is a smaller population with a lower probability of default. Performance metrics drop marginally relative to those for the model applied to the population of all borrowers but they are still very strong. For example, the AUC score drops from 92-93% to 86-88%, accuracy mostly remain in the same range and the loss increases by 1-2 percentage points.

## 5.2 Model Interpretation

We use our hybrid DNN-GBT model to uncover associations between the explanatory variables and default behavior. Since we do not identify causal relationships, our goal is simply to find covariates that have an important impact on default outcomes. Our findings can be used to better

Table 5: Performance Metrics using Hybrid DNN-GBT

Training Window	Testing Window	AUC score	Precision	Recall	F-measure	Accuracy	Loss
2004Q1-2013Q4	2004Q1-2013Q4	0.9554	0.8703	0.8160	0.8470	0.8962	0.2535
2004Q1	2006Q1	0.9246	0.8531	0.7016	0.7700	0.8638	0.3233
2004Q2	2006Q2	0.9255	0.8509	0.7151	0.7771	0.8657	0.3178
2004Q3	2006Q3	0.9264	0.8488	0.7273	0.7834	0.8670	0.3156
2004Q4	2006Q4	0.9255	0.8499	0.7258	0.7829	0.8653	0.3187
2005Q1	2007Q1	0.9260	0.8599	0.7191	0.7832	0.8643	0.3205
2005Q2	2007Q2	0.9260	0.8617	0.7202	0.7846	0.8638	0.3211
2005Q3	2007Q3	0.9253	0.8606	0.7211	0.7847	0.8627	0.3246
2005Q4	2007Q4	0.9241	0.8554	0.7280	0.7866	0.8615	0.3276
2006Q1	2008Q1	0.9251	0.8513	0.7390	0.7912	0.8621	0.3262
2006Q2	2008Q2	0.9249	0.8569	0.7295	0.7881	0.8609	0.3271
2006Q3	2008Q3	0.9259	0.8563	0.7347	0.7909	0.8618	0.3248
2006Q4	2008Q4	0.9263	0.8582	0.7347	0.7916	0.8613	0.3247
2007Q1	2009Q1	0.9279	0.8512	0.7509	0.7979	0.8635	0.3202
2007Q2	2009Q2	0.9283	0.8513	0.7538	0.7996	0.8647	0.3193
2007Q3	2009Q3	0.9289	0.8439	0.7665	0.8034	0.8659	0.3166
2007Q4	2009Q4	0.9306	0.8474	0.7710	0.8074	0.8680	0.3124
2008Q1	2010Q1	0.9304	0.8462	0.7721	0.8074	0.8678	0.3137
2008Q2	2010Q2	0.9301	0.8367	0.7820	0.8084	0.8678	0.3143
2008Q3	2010Q3	0.9298	0.8340	0.7849	0.8087	0.8678	0.3144
2008Q4	2010Q4	0.9295	0.8326	0.7823	0.8066	0.8671	0.3157
2009Q1	2011Q1	0.9304	0.8379	0.7806	0.8082	0.8689	0.3133
2009Q2	2011Q2	0.9289	0.8329	0.7796	0.8054	0.8677	0.3161
2009Q3	2011Q3	0.9294	0.8416	0.7694	0.8039	0.8686	0.3145
2009Q4	2011Q4	0.9296	0.8354	0.7756	0.8044	0.8686	0.3135
2010Q1	2012Q1	0.9302	0.8332	0.7785	0.8049	0.8692	0.3115
2010Q2	2012Q2	0.9291	0.8319	0.7718	0.8007	0.8681	0.3134
2010Q3	2012Q3	0.9289	0.8278	0.7735	0.7997	0.8684	0.3119
2010Q4	2012Q4	0.9282	0.8199	0.7749	0.7968	0.8673	0.3136
2011Q1	2013Q1	0.9288	0.8156	0.7800	0.7974	0.8676	0.3121
2011Q2	2013Q2	0.9284	0.8169	0.7758	0.7958	0.8679	0.3117
2011Q3	2013Q3	0.9290	0.8101	0.7834	0.7965	0.8680	0.3103
2011Q4	2013Q4	0.9293	0.8108	0.7805	0.7954	0.8685	0.3083

Performance metrics for our model of default risk. The model calibrations are specified by the training and testing windows. The results of classifications versus actual outcomes over the following 8Q are used to calculate these performance metrics for 90+ days delinquencies within 8Q. Source: Authors' calculations based on Experian Data.

Table 6: Performance Metrics using Hybrid DNN-GBT

Training Window	Testing Window	AUC score	Precision	Recall	F-measure	Accuracy	Loss
2004Q1	2006Q1	0.8776	0.7688	0.4035	0.5292	0.8676	0.3167
2004Q2	2006Q2	0.8658	0.7315	0.3543	0.4774	0.8680	0.3159
2004Q3	2006Q3	0.8644	0.7189	0.3647	0.4839	0.8681	0.3159
2004Q4	2006Q4	0.8649	0.7236	0.3640	0.4843	0.8661	0.3189
2005Q1	2007Q1	0.8648	0.7425	0.3571	0.4823	0.8616	0.3284
2005Q2	2007Q2	0.8624	0.7364	0.3477	0.4724	0.8592	0.3320
2005Q3	2007Q3	0.8613	0.7376	0.3457	0.4708	0.8577	0.3368
2005Q4	2007Q4	0.8606	0.7273	0.3579	0.4797	0.8566	0.3386
2006Q1	2008Q1	0.8621	0.7206	0.3815	0.4989	0.8551	0.3405
2006Q2	2008Q2	0.8618	0.7267	0.3631	0.4842	0.8533	0.3420
2006Q3	2008Q3	0.8612	0.7194	0.3627	0.4823	0.8542	0.3401
2006Q4	2008Q4	0.8583	0.7130	0.3484	0.4681	0.8562	0.3363
2007Q1	2009Q1	0.8599	0.7001	0.3771	0.4902	0.8603	0.3285
2007Q2	2009Q2	0.8600	0.6978	0.3786	0.4909	0.8624	0.3251
2007Q3	2009Q3	0.8599	0.6802	0.4002	0.5039	0.8650	0.3204
2007Q4	2009Q4	0.8581	0.6740	0.3905	0.4945	0.8674	0.3162
2008Q1	2010Q1	0.8613	0.6858	0.4065	0.5105	0.8687	0.3146
2008Q2	2010Q2	0.8622	0.6687	0.4334	0.5259	0.8696	0.3129
2008Q3	2010Q3	0.8622	0.6622	0.4454	0.5326	0.8701	0.3125
2008Q4	2010Q4	0.8637	0.6675	0.4384	0.5292	0.8718	0.3091
2009Q1	2011Q1	0.8669	0.6785	0.4517	0.5424	0.8724	0.3084
2009Q2	2011Q2	0.8670	0.6770	0.4538	0.5434	0.8728	0.3085
2009Q3	2011Q3	0.8685	0.6923	0.4341	0.5336	0.8734	0.3069
2009Q4	2011Q4	0.8657	0.6773	0.4278	0.5244	0.8760	0.3027
2010Q1	2012Q1	0.8672	0.6774	0.4289	0.5253	0.8756	0.3019
2010Q2	2012Q2	0.8693	0.6895	0.4234	0.5247	0.8755	0.3015
2010Q3	2012Q3	0.8682	0.6839	0.4216	0.5217	0.8764	0.3000
2010Q4	2012Q4	0.8676	0.6811	0.4133	0.5144	0.8771	0.2984
2011Q1	2013Q1	0.8700	0.6754	0.4450	0.5365	0.8765	0.2996
2011Q2	2013Q2	0.8705	0.6789	0.4401	0.5340	0.8769	0.2986
2011Q3	2013Q3	0.8698	0.6711	0.4424	0.5333	0.8778	0.2967
2011Q4	2013Q4	0.8681	0.6681	0.4317	0.5245	0.8789	0.2950

Performance metrics for our model of default risk for the current population. Borrowers who are current do not have any delinquencies. The model calibrations are specified by the training and testing windows. The results of classifications versus actual outcomes over the following 8Q are used to calculate these performance metrics for 90+ days delinquencies within 8Q. Source: Authors' calculations based on Experian Data.

understand default behavior, further refine model specification and possibly aid in the formulation of theoretical models of consumer default. For this exercise, we mainly use the pooled model, which uses all available data. This allows us to assess factors that are critical in default behavior throughout the sample period with the best performing model. We also consider time variation in the factors influencing the default decision in subsets of our sample.

### 5.2.1 Explanatory Power of Variables

We start by examining the explanatory power of each of our features. We follow an approach similar to Sirignano, Sadhwani, and Giesecke (2018), which amounts to a perturbation analysis on the pooled sample using our hybrid model. First, we draw a random sample of 100,000 observations from the testing sample. Then, for each variable, we re-shuffle the feature, keeping the distribution intact and the model’s loss function is evaluated with the changed covariate. We repeat this step 10 times, and report the average of the loss and accuracy. Then, the variable is replaced to its original values, and a perturbation test is performed on a new variable. Perturbing the variable of course reduces the accuracy of the model, and the test loss becomes larger. If a particular variable has strong explanatory power, the test loss will significantly increase. The test loss for the complete model when no variables are perturbed is the Baseline value. Features that have large explanatory power, and whose information is not contained in the other remaining variables will increase the loss significantly if they are altered. Table 7 reports the results. Features relating to the number and credit available on revolving trades and features pertaining to delinquencies dominate the list. Specifically, credit amount on revolving trades, months since the most recent 90+ days delinquency and months since the oldest trade was opened each increase the loss by 15%, the number of open bankcard and credit card trades increase the loss by 13%, while the balance on first mortgage trades and monthly payment on first mortgages increase the loss by 11%. These results suggest that revolving debt, length of credit history and temporal proximity to a delinquency are all important factors in default behavior. Based on publicly available information, length of the credit history is also an important determinant of standard credit scoring models, though credit utilization rather than balances or number of trades is understood as the most critical. This approach to assessing the importance of different features for the predicted probability of default has two major shortcomings. First, when features are highly correlated, the interpretation of feature importance can be biased by unrealistic data instances. To illustrate this problem, consider two highly correlated features. As we perturb one of the features, we create instances that are unlikely or even impossible. For example, mortgage balances are highly correlated with and lower than total debt balances, yet this perturbation approach could create instances in which total debt balances are smaller than mortgage balances. Since many of the features are strongly correlated, care must be taken with interpretation of feature importance. We list the highly correlated features in Appendix C. An additional concern with this perturbation approach is that the distribution of some features are highly skewed, which implies that the probability of their value being different than where the mass

of their distribution is concentrated is quite low. Moreover, skewness varies substantially across features, therefore the informativeness of the perturbation may differ across variables. In the next section, we examine a more robust approach that is less susceptible to these limitations.

### 5.2.2 Economic Significance of Variables

We now turn to analyzing the economic significance of our features for default behavior. We adopt SHapley Additive exPlanations (SHAP), a unified framework for interpreting predictions, to explain the output of our hybrid deep learning model (for a detailed description of the approach see Lundberg and Lee (2017)). SHAP uses a game theoretical concept to assign each feature a local importance value for a given prediction.

To obtain the Shapley values for our hybrid model, we first compute the Shapley values for the Deep Neural Network model and the Gradient Boosted Trees model separately, then simply take the average of the Shapley values for these two models for each individual and for each feature.<sup>15</sup> We use a random sample of 1 million observations for explaining the model. Features that are highly correlated can decrease the importance of the associated feature by splitting the importance between both features. To account for the effect of feature correlation on interpretability, we group features with a correlation larger than 0.7, and sum the SHAP values within each groups. We denote these groups with an asterisk for the rest of the analysis and report the composition of feature groups in Table 19 in the appendix.

Figure 7 sorts features by the sum of SHAP value magnitudes, and plots the distribution of the impact each feature has on the model output for the twelve most important features or groups of correlated features. The color represents the feature value (red: high, blue: low), whereas the position on the horizontal axis denotes whether the feature is above or below the reference value and by how much. The charts plot the distribution of SHAP values for individual instances in the 1 million testing sample. The most important feature in terms of SHAP value magnitude is the worst status on any trades. High values of this variable tend to be highly associated with a higher value of the default rate compared to the baseline or reference value, whereas low values to a lower than baseline default rate, though the distribution of instances is much more dispersed. Features capturing credit history, such as length of credit history and recent delinquencies, also have high SHAP values, specifically, higher than baseline values of this variable are associated mostly with lower than baseline values of the predicted default probability, while lower than baseline values of the credit history are associated with a higher than baseline values of the default rate, with a much more dispersed distribution. Additionally, credit card utilization, credit amount on derogatory trades and outstanding collections are typically associated with an increase in the predicted probability of

---

<sup>15</sup>We implement Deep SHAP, a high-speed approximation algorithm for SHAP values in deep learning models to compute the Shapley values for our 5 hidden layer neural network. Because our dataset is fairly large with many features, we do not use the entire training dataset for the prediction. We pass a random sample of 1,000 observations of the training dataset. For GBT, we implement TreeExplainer, a high-speed exact algorithm for tree ensemble methods. For this, we do not need background or training observations.

Table 7: Explanatory Power of Variables

Feature	Accuracy	Loss
Credit amount on revolving trades	0.8792	0.2939
Months since the most recent 90+ days delinquency	0.8795	0.2938
Months since the oldest trade was opened	0.8770	0.2932
Open credit card trades	0.8793	0.2889
Open bankcard revolving, and charge trades	0.8803	0.2854
Balance on first mortgage trades	0.8772	0.2829
Monthly payment on open first mortgage trades	0.8752	0.2821
Total credit amount on open trades	0.8830	0.2821
Credit amount on open credit card trades	0.8829	0.2815
Worst ever status on any trades in the last 24 months	0.8869	0.2805
Monthly payment on all debt	0.8825	0.2803
Total debt balances	0.8846	0.2801
Balance on bankcard revolving and charge trades	0.8860	0.2747
Balance on collections	0.8842	0.2743
Months since the most recently opened first mortgage	0.8856	0.2743
Monthly payment on open auto loan trades	0.8857	0.2739
Monthly payment on credit card trades	0.8848	0.2733
Balance on credit card trades	0.8867	0.2732
Months since the most recent 30-180 days delinquency	0.8865	0.2731
Credit amount on open mortgage trades	0.8868	0.2724
Credit amount on unsatisfied derogatory trades	0.8848	0.2717
Balance on revolving trades	0.8860	0.2717
Credit amount on open installment trades	0.8869	0.2713
Months since the most recently opened credit card trade	0.8869	0.2709
Worst present status on any trades	0.8888	0.2706
Months since the most recently opened auto loan trade	0.8870	0.2700
Early payoff trades	0.8879	0.2695
Credit amount paid down on open first mortgage trades	0.8898	0.2685
Mortgage type	0.8885	0.2679
Months since the most recently closed, transferred, or refinanced first mortgage	0.8887	0.2671
Fraction of mortgage to total debt	0.8896	0.2669
Credit card utilization ratio	0.8881	0.2665
Balance on open auto loan trades	0.8887	0.2663
Credit amount on non-deferred student trades	0.8901	0.2662
:	:	:
:	:	:
Mortgage inquiries made in the last 3 months	0.8948	0.2552
Bankcard revolving and charge inquiries made in the last 3 months	0.8947	0.2551
Auto loan or lease inquiries made in the last 3 months	0.8946	0.2551
Balance on open bankcard revolving, and charge trades with credit line suspended	0.8947	0.2551
Baseline	0.8947	0.2551

This table reports a perturbation analysis on the pooled sample using our hybrid model. For each variable, we re-shuffle the feature, keeping the distribution intact in the test dataset and the model’s loss function is evaluated on the test dataset with the changed covariate. We repeat this step 10 times, and report the average of the loss and accuracy. Then, the variable is replaced to its original values, and a perturbation test is performed on a new variable. Perturbing the variable of course reduces the accuracy of the model, and the test loss becomes larger. If a particular variable has strong explanatory power, the test loss will significantly increase. The test loss for the complete model when no variables are perturbed is the Baseline value.

default relative to the baseline, as are relatively high numbers of inquiries and monthly payment on debt. Higher total debt balances are also associated with a lower than baseline predicted probability of default, reiterating the notion that the borrowers with the most credit are also associated with lower predicted probability of default, which suggests that credit allocation decisions are made to minimize default probabilities. As in the perturbation exercise, we find that number of trades and balances seem to have the strongest association with variation in the predicted probability of default.

These results only point to correlations between the features and the predicted outcome and should not be interpreted causally. Yet, they can be used as a point of departure for a causal analysis of default and theoretical modeling. They are also important to comply with legal disclosure requirements. Both the Fair Credit Reporting Act and the Equal Opportunity in Credit Access Act require lenders and developers of credit scoring models to reveal the most important factors leading to a denial of a credit application and for credit scores. The SHAP value provides an individualized assessment of such factors that can be used for making credit allocation decisions and communicating them to the borrower.




Figure 7: SHAP applied to predicted 90+ days delinquency within 8Q. Source: Authors’ calculations based on Experian Data.

### 5.2.3 Temporal Determinants of Default

We next look at the changing dynamics of default behavior by comparing models that are trained in different periods of time. For this analysis, we use our hybrid model. Specifically, we target

the following time periods: 2006Q1, 2008Q1 and 2011Q1 as time periods before, during and after the 2007-2009 crisis, and compute default predictions for them with data trained in the same quarter and two years prior, that is in 2004Q1, 2006Q1, 2007Q1 and 2009Q1, respectively. We then calculate Shapley values for the two models.<sup>16</sup> The first exercise provides an in-sample assessment for feature importance, while the second exercise can be used to assess feature importance out-of-sample. In both exercises, the model is the same, so comparing the results from the two exercises can help uncover which features are important for default prediction for a given period from an ex ante perspective and from an in-sample perspective.

Table 8 reports the results. The features are sorted by the sum of absolute SHAP value magnitudes over the first period. For each period, it is interesting to compare the variation in SHAP values from an ex ante and contemporaneous perspective, and additionally we are interested in comparing variation in SHAP values for given features in the different time periods. In all testing windows, the worst status on a mortgage trade has the first or second highest SHAP value. Total debt balances are the feature with the first or second highest SHAP value in all testing periods except for 2008Q1 where for the within same quarter prediction they are only the fifth in rank. Student debt is the third ranked feature in terms of SHAP values in all testing periods, and trades every more 90+ days delinquent or derogatory is either fourth or fifth. In 2006Q1 and 2011Q1, number of open credit card trades come in fifth, but ranks lower in 2008Q1, while features pertaining to credit card debt rank in a range for fourth to seventh depending on the period. Fraction of 60 days delinquent debt to total debt ranks sixth to tenth depending on the time period. The SHAP value is quite stable over time for most features, but there are some variables for which it changes substantially. One example is foreclosed first mortgages which ranks below fifteenth for 2006Q1 and 2008Q1 but moves up to eighth out of sample and ninth in sample for 2008Q1. The length of the credit history does not register high SHAP value in any of the time periods, nor do features related to credit utilization or inquiries. Overall these results confirm our findings from the pooled model, suggesting the balances and number of trades, in addition to delinquency status, have a strong association with default risk according to our model.

### 5.3 Model Comparisons

We examine the out-of-sample behavior of a collection of alternative machine learning techniques in this section. To motivate our choice of deep learning, leading to our hybrid DNN-GBT model, we begin by illustrating the importance of hidden layers that enable us to capture multi-level interactions between features by comparing how neural networks of different depth perform on the pooled sample. For this exercise, we fix the number of neurons per layer at 512, and build neural

---

<sup>16</sup>We do this for both the Deep Neural Network and the Gradient Boosted Trees and similarly to how we obtain the output, we simply take the average of the Shapley values. For our DNN explanations, we use a random sample of 1,000 observations of the testing data scaled by the mean and standard deviation of the corresponding testing data for reference value.



Table 8: SHAP Values over Time

Features	Prediction Date					
	2006Q1		2008Q1		2011Q1	
	2004Q1	2006Q1	2006Q1	2008Q1	2009Q1	2011Q1
Worst present status on a mortgage trade	0.446 (1)	0.485 (2)	0.48 (2)	0.439 (1)	0.472 (1)	0.395 (2)
Total debt balances*	0.301 (2)	0.539 (1)	0.497 (1)	0.206 (5)	0.416 (2)	0.473 (1)
Student debt*	0.282 (3)	0.271 (3)	0.242 (3)	0.247 (3)	0.302 (3)	0.317 (3)
Trades ever 90 or more days delinquent or derogatory*	0.22 (4)	0.236 (5)	0.235 (4)	0.237 (4)	0.283 (4)	0.245 (4)
Number of open credit cards*	0.215 (5)	0.038 (19)	0.033 (22)	0.16 (7)	0.042 (20)	0.202 (5)
Fraction of 60 days delinquent debt to total debt	0.122 (6)	0.185 (6)	0.186 (6)	0.119 (8)	0.075 (10)	0.099 (10)
Credit card debt*	0.105 (7)	0.26 (4)	0.216 (5)	0.18 (6)	0.09 (7)	0.163 (6)
Months since the most recently opened home equity line of credit trade	0.089 (8)	0.036 (21)	0.035 (21)	0.08 (11)	0.06 (14)	0.148 (8)
Credit amount on unsatisfied derogatory trades	0.085 (9)	0.106 (8)	0.087 (9)	0.1 (10)	0.094 (6)	0.078 (13)
Balance on HELOC loans*	0.077 (10)	0.179 (7)	0.171 (7)	0.023 (32)	0.016 (41)	0.088 (11)
Months since the most recent 30-180 days delinquency	0.076 (12)	0.075 (11)	0.076 (11)	0.039 (21)	0.044 (19)	0.064 (16)
Auto loan*	0.076 (11)	0.037 (20)	0.04 (20)	0.267 (2)	0.098 (5)	0.036 (23)
Worst status on any trades*	0.071 (13)	0.091 (9)	0.09 (8)	0.063 (15)	0.058 (15)	0.052 (19)
90+ days late debt balances*	0.062 (14)	0.055 (15)	0.061 (13)	0.023 (33)	0.029 (26)	0.045 (20)
Foreclosed first mortgages*	0.048 (15)	0.04 (18)	0.041 (19)	0.033 (24)	0.082 (8)	0.108 (9)
Balance on open 60 days late revolving trades	0.044 (16)	0.06 (13)	0.058 (14)	0.038 (22)	0.038 (23)	0.088 (12)
Number of mortgages*	0.043 (17)	0.065 (12)	0.067 (12)	0.05 (17)	0.052 (17)	0.074 (14)
Fraction of credit card debt to total debt	0.034 (18)	0.015 (38)	0.016 (37)	0.026 (30)	0.032 (25)	0.054 (17)
Charge-off amount on unsatisfied charge-off trades	0.031 (19)	0.058 (14)	0.051 (16)	0.04 (20)	0.007 (68)	0.033 (25)
Home equity line of credit trades ever 90 or more days delinquent or derogatory	0.028 (20)	0.018 (33)	0.016 (36)	0.013 (49)	0.014 (47)	0.027 (28)

This table reports the Shapley values for a selected 20 features for three out-of-sample models. For each prediction window, we compute the Shapley value for each of the observations and for each feature. We then calculate the average of the absolute value for each feature, and report the results for the selected features. Finally, we rank the results based on the feature's relative rank in the given prediction window in parentheses. Source: Authors' calculations based on Experian data.

networks up to 8 hidden layers.<sup>17</sup>

We benchmark our results against the logistic regression, which is a commonly used technique in credit scoring and can be interpreted as a neural network with no hidden layers. Table 9 reports the in- and out-of-sample behavior for neural networks with 0-8 hidden layers. The number of hidden layers measure the complexity of the network, and we found that the marginal improvements in performance beyond 3 layers are small. Depth, however, is not directly proportional to out-of-sample performance due to higher model capacity. Table 9 also shows that applying dropout substantially improves the out-of-sample fit for networks of higher depths. This demonstrates that dropout serves as an effective regularization tool and addresses over-fitting for networks of greater depths.

Table 9: Neural networks comparison: Loss & Accuracy

Model	In-sample Loss		Out-of-sample Loss	
	w/o Dropout	Dropout	w/o Dropout	Dropout
Logistic Regression	0.3451	0.3452	0.3444	0.3446
1 layer	0.3097	0.3113	0.3123	0.3116
2 layers	0.2987	0.2983	0.3089	0.3041
3 layers	0.2878	0.2888	0.3074	0.2987
4 layers	0.2850	0.2862	0.3078	0.2972
5 layers	0.2856	0.2867	0.3076	0.2975
6 layers	0.2850	0.2852	0.3075	0.2967
7 layers	0.2825	0.2838	0.3066	0.2958
8 layers	0.2901	0.2853	0.3062	0.2966

Model	In-sample Accuracy		Out-of-sample Accuracy	
	w/o Dropout	Dropout	w/o Dropout	Dropout
Logistic Regression	0.8564	0.8564	0.8569	0.8570
1 layer	0.8693	0.8684	0.8683	0.8685
2 layers	0.8745	0.8748	0.8701	0.8722
3 layers	0.8798	0.8789	0.8714	0.8742
4 layers	0.8811	0.8799	0.8715	0.8747
5 layers	0.8807	0.8801	0.8721	0.8747
6 layers	0.8813	0.8803	0.8721	0.8749
7 layers	0.8821	0.8809	0.8723	0.8753
8 layers	0.8785	0.8807	0.8719	0.8750

In-sample and out-of-sample loss (categorical cross-entropy) and accuracy for neural networks of different depth and for logistic regression. Models are calibrated and evaluated on the pooled sample (2004Q1 - 2013Q4). Source: Authors' calculations based on Experian Data.

The results in Table 9 suggest that there are complex non-linear relationships among the fea-

<sup>17</sup>The architecture reported in Table 9 was not optimized. We picked 512 for each layer randomly. Another variant of this exercise is where we use the optimal 5-layer neural network architecture and we remove the neurons layer by layer. We report the results in Table 23.

tures used as inputs in the model. This is further supported by the fact that permitting non-linear relationships between default behavior and explanatory variables produces the largest model improvement. Going from a linear model (0 layers) to the simplest non-linear model (1 layer) generates the most sizable reduction in out-of-sample loss. To see this from another angle, we plot the ROC curves for our neural networks considered in Figure 8. We can see that the logistic regression is dominated by all models that allow for non-linear relationships, while the improvements for deeper models are marginal.




Figure 8: Out-of-sample ROC curves for various models. Models are calibrated and evaluated on the pooled sample (2004Q1 - 2013Q4). Source: Authors' calculations based on Experian Data.

We next analyze a number of machine learning techniques that are possible alternatives to our hybrid model. These algorithms have been used in other credit scoring applications, and include decision trees (CART, see Khandani, Kim, and Lo (2010)), random forests (RF, see Butaru et al. (2016)), neural networks (see West (2000)), gradient boosting (GBT, see Xia et al. (2017)) and logistic regression. We use the out of sample loss as our main comparison metric, with lower loss values corresponding to better model performance. We tune the hyper-parameters for each model and present the results in Table 10 for our baseline 1 quarter training/validation samples. Our hybrid model performs the best, with gradient boosting coming second, while deep neural network outperforms both random forest and decision trees in all samples. We repeat this comparison with an expanding training window and find that deep neural networks benefit the most from more training data.

Empirically, ensembles perform better when there is a significant diversity among the models

(see Kucheva and Whitaker (2003)). Table 22 in Appendix D shows the SHAP values for our hybrid DNN-GBT model in comparison to GBT and DNN models for the pooled sample. The results suggest there are significant differences between the DNN and GBT. For instance, total debt balances is the third most important feature for GBT, while only tenth for DNN. Even more striking perhaps is that monthly payment on all debt is ranked as tenth most important for GBT, while only 45th for DNN. The ensemble approach can thus be thought of as providing diversification, which can reduce the variance of any of the two models. If one of the models puts too high of a weight on a feature, the other model may mitigate this effect.

It is important to emphasize that these results do not imply that there does not exist a random forest or CART model that cannot outperform our hybrid model. The best model will depend on the specific sample. The exercise is intended to illustrate that the complexity of the model is proportional to its accuracy to a certain degree, and that deep neural networks improve substantially on shallow models, such as logistic regression.

## 6 Applications

In this section, we use our model in a number of applications. We first provide a comparison between the performance of our model and conventional credit score. Second, we show that our model can score more borrowers. Finally, we show that our model is able to accurately predict variations in aggregate default risk.

### 6.1 Comparison with Credit Score

In this section, we compare the performance of our deep neural networks to a conventional credit score.<sup>18</sup> The credit score is a summary indicator intended to predict the risk of default by the borrower and it is widely used by the financial industry. For most unsecured debt, lenders typically verify a prospective borrower’s credit score at the time of application and sometimes a short recent sample of their credit history. For larger unsecured debts, lenders also typically require some form of income verification, as they do for secured debts, such as mortgages and auto loans. Still, the credit score is often a key determinant of crucial terms of the borrowing contract, such as the interest rate, the downpayment or the credit limit. We have access to a widely used conventional credit score that uses information from the three credit bureaus.

Table 11 shows the relationship between credit score, predicted probability and realized default rate, where default is defined as usual as 90+ days delinquency in the subsequent 8 quarters. The calculation proceeds as follows. We first compute the number of unique credit scores in the data. We create the same number of bins of equal size in our predicted probability distribution, and calculate the realized frequency of 90+ days delinquencies in the subsequent 8 quarters for each

---

<sup>18</sup>The deep-learning forecasts for each quarter are obtained using out-of-sample input data, as reported in Table 3.

Table 10: Model Comparison: Out-of-Sample Loss

Training Window	Testing Window	Combined	GBT	DNN	RF	CART	Logistic
2004Q1	2006Q1	0.3233	0.3235	0.3276	0.3274	0.3432	0.3499
2004Q2	2006Q2	0.3178	0.3185	0.3220	0.3231	0.3372	0.3465
2004Q3	2006Q3	0.3156	0.3161	0.3202	0.3210	0.3361	0.3445
2004Q4	2006Q4	0.3187	0.3198	0.3223	0.3245	0.3396	0.3464
2005Q1	2007Q1	0.3205	0.3209	0.3249	0.3259	0.3427	0.3491
2005Q2	2007Q2	0.3211	0.3215	0.3252	0.3272	0.3440	0.3512
2005Q3	2007Q3	0.3246	0.3246	0.3294	0.3301	0.3455	0.3525
2005Q4	2007Q4	0.3276	0.3275	0.3323	0.3327	0.3498	0.3560
2006Q1	2008Q1	0.3262	0.3264	0.3308	0.3316	0.3477	0.3558
2006Q2	2008Q2	0.3271	0.3277	0.3316	0.3321	0.3495	0.3563
2006Q3	2008Q3	0.3248	0.3256	0.3285	0.3299	0.3474	0.3565
2006Q4	2008Q4	0.3247	0.3259	0.3282	0.3293	0.3462	0.3547
2007Q1	2009Q1	0.3202	0.3211	0.3248	0.3244	0.3425	0.3503
2007Q2	2009Q2	0.3193	0.3202	0.3241	0.3236	0.3427	0.3496
2007Q3	2009Q3	0.3166	0.3174	0.3213	0.3205	0.3383	0.3464
2007Q4	2009Q4	0.3124	0.3127	0.3172	0.3170	0.3348	0.3425
2008Q1	2010Q1	0.3137	0.3141	0.3185	0.3182	0.3360	0.3453
2008Q2	2010Q2	0.3143	0.3144	0.3193	0.3188	0.3355	0.3463
2008Q3	2010Q3	0.3144	0.3148	0.3195	0.3187	0.3353	0.3468
2008Q4	2010Q4	0.3157	0.3156	0.3214	0.3196	0.3365	0.3459
2009Q1	2011Q1	0.3133	0.3138	0.3179	0.3174	0.3320	0.3453
2009Q2	2011Q2	0.3161	0.3163	0.3211	0.3202	0.3346	0.3472
2009Q3	2011Q3	0.3145	0.3148	0.3192	0.3186	0.3322	0.3457
2009Q4	2011Q4	0.3135	0.3137	0.3181	0.3176	0.3313	0.3446
2010Q1	2012Q1	0.3115	0.3118	0.3160	0.3166	0.3298	0.3440
2010Q2	2012Q2	0.3134	0.3140	0.3177	0.3186	0.3313	0.3453
2010Q3	2012Q3	0.3119	0.3127	0.3160	0.3176	0.3310	0.3440
2010Q4	2012Q4	0.3136	0.3142	0.3180	0.3195	0.3325	0.3455
2011Q1	2013Q1	0.3121	0.3124	0.3168	0.3167	0.3313	0.3440
2011Q2	2013Q2	0.3117	0.3123	0.3158	0.3169	0.3320	0.3432
2011Q3	2013Q3	0.3103	0.3108	0.3148	0.3152	0.3315	0.3416
2011Q4	2013Q4	0.3083	0.3088	0.3126	0.3136	0.3281	0.3406
Average		0.3171	0.3177	0.3216	0.3220	0.3377	0.3476

Performance comparison of machine learning classification models of consumer default risk. The model calibrations are specified by the training and testing windows. The results of predicted probabilities versus actual outcomes over the following 8Q testing period are used to calculate the loss metric for 90+ days delinquencies within 8Q. Combined refers to the hybrid DNN-GBT model, DNN refers to deep neural network, RF refers to random forest, GBT refers to gradient boosted trees, while CART refers to decision tree. Source: Authors' calculations based on Experian Data.

of these bins. Since higher credit scores correspond to lower probability of default, we present the negative of the rank correlation with realized defaults for the credit score. The results indicate that even though credit score is successful in rank-ordering customers by their future default rates, with rank correlations between 0.983 and 0.990, our deep neural network performs better, with rank correlations always at 0.999. Figure 14 in Appendix E plots the time series of these rank correlations by quarter for the entire sample period. The figure shows that the rank correlation for the predicted probability of default generated by our model is remarkably stable over time, while for the credit score it fluctuates from lows of around 0.975 before 2012 to a peak of 0.995 in 2013Q2 with notable quarter by quarter variation. This property of credit scores may be due to the fact that the credit score is an ordinal ranking and its distribution is designed to be stable over time, even if default risk at an individual or aggregate level may change substantially. Figure 13 in Appendix E displays the histogram of credit score distributions in our sample for selected years, and show that these distributions are virtually identical over time.

A common way to measure the accuracy of conventional credit scoring models is the Gini coefficient, which measures the dispersion of the credit score distribution and therefore its ability to separate borrowers by their default risk. The Gini coefficient is related to a key performance metric for machine learning algorithm, the AUC score, with  $Gini = 2 * AUC - 1$ , so we can compare the performance of the credit score to our model along this dimension. Figure 14 plots the Gini coefficient for the credit score and our predicted default probability by quarter. The Gini coefficient for our model is about 0.85 between 2006Q1 and 2008Q3, and then rises to 0.86. For the credit score, the Gini coefficient is close to 0.81 until 2012Q3 when it drops to approximately 0.79 until the end of the sample, suggesting a drop in performance of the credit score in the aftermath of the Great Recession.

Table 11: Credit Score-Predicted Probability of Default Comparison

Metric	2006	2007	2008	2009	2010	2011	2012	2013
<b>Rank Correlation</b>								
Credit Score	0.9881	0.9804	0.9882	0.9816	0.9825	0.9861	0.9906	0.9944
Predicted Probability	0.9991	0.9992	0.9992	0.9992	0.9992	0.9991	0.9993	0.9992
<b>GINI Coefficient</b>								
Credit Score	0.8108	0.8142	0.8137	0.8143	0.8078	0.8008	0.7942	0.7898
Predicted Probability	0.8538	0.8535	0.8536	0.8606	0.8616	0.8608	0.8601	0.8580

Rank correlation between credit score, predicted probability of default according to our model and subsequent realized default frequency by year. For the credit score, report the rank correlation between each unique value of the score and the default frequency. For predicted probability of default based on our hybrid DNN-GBT model, we first generate a number of bins equal to the number of unique credit score realizations in the data and then calculate the realized default frequency for each bin. Source: Authors' calculations based on Experian Data.

Figure 9 plots a scatterplot of the realized default rate against the credit score (left panel) and

our predicted probability (right panel) for all quarters in the year 2008. In addition to the raw data, we also plot second-order polynomial-fitted curves to approximate the relationship. The scatter plots of realized default rates against the predictions from our hybrid model lay mostly on the 45 degree line, consistent with the very high rank correlations reported in Table 11. By contrast, the relation between realized default rates and credit scores has an inverted S-shape, with the realized default rate equal to one for a large range of low credit scores and equal to zero for a large range of high credit scores, and a large variation only for intermediate credit scores.




Figure 9: Scatter plot of realized default rates against model predicted default probability (a) and the credit score (b), with associated second-order polynomial fitted approximations for the year 2008. Source: Authors’ calculations based on Experian Data.

Figure 10 plots second-order polynomial-fitted curves approximating the relation between realized default rates and those predicted by our model and the credit score for all years in which our model prediction is available, starting in 2006 until 2013, to examine how the relation between realized and predicted defaults varies with aggregate economic conditions. For the years at the height of the Great Recession, the default rate seems to be somewhat higher than our model prediction, but in all years the relation is very close to a 45 degree line. By contrast, there is virtually no change in the relation between the realized default rate and the credit score. This is by construction, since the distribution of credit scores is designed to only provide a relative ranking of default risk across borrowers.<sup>19</sup> This property of the credit score implies that it is unable to forecast variations in aggregate default risk. In Section 6.2, we will show that our model is able to capture variations in aggregate default risk while retaining a consistent ability to separate borrowers by their individual default risk.

Table 12 reports the rank correlation with the realized default rate and the Gini coefficients by

<sup>19</sup>Credit scores are specifically designed to provide a stable ranking by using multiple years of data.

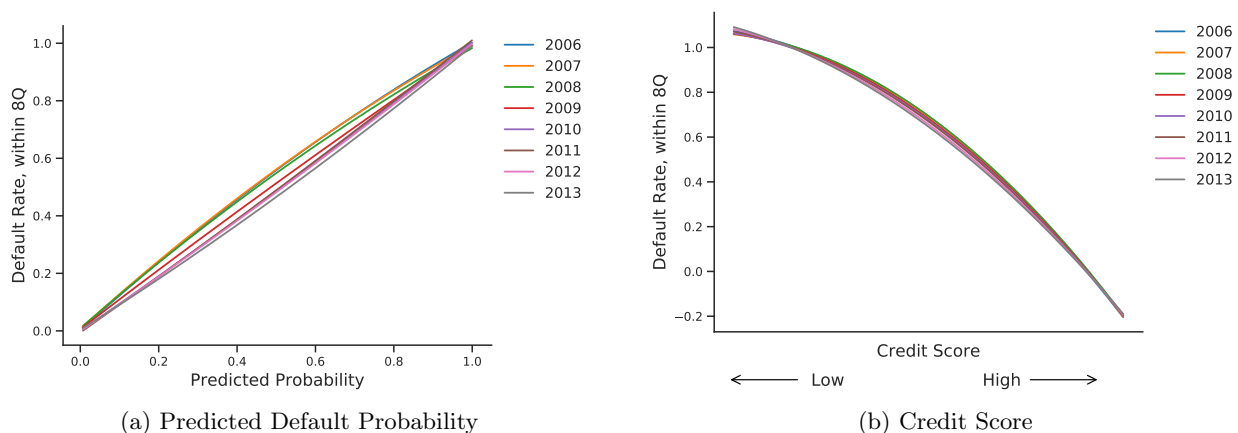


Figure 10: Second-order polynomial approximation of the relationship between realized default rates against model predicted default probability (a) and the credit score (b) for selected years. Source: Authors’ calculations based on Experian Data.

year for the credit score and the probability of default predicted by our model restricting attention to the current population, that is those borrowers who do not have any outstanding delinquencies in the quarter of interest. The rank correlation between the credit score and the realized default rate drops by 1-3 percentage points for these borrowers, whereas for our model it drops by less than a quarter of 1 percent. The Gini coefficient drops from 80-81% to 68-69% for the credit score and from 85-86% to 72-74%. These results suggest that when measured on the population of current borrowers, the performance advantage of our model relative to a conventional credit score grows.<sup>20</sup>

To understand the differences in performance across the credit score and our model, we examine how the ranking of borrowers varies under the two approaches. To do so, we consider the industry classification of borrowers into five risk categories Deep Subprime, Subprime, Near Prime, Prime and Super Prime.<sup>21</sup> As shown in Table 13, these categories account for respectively, 6.5%, 21.2%, 14.1%, 33.3% and 24.9% of all borrowers. We then create 5 correspondingly sized bins in our predicted probability of default with bin 1 corresponding to the 6.5% of borrowers with the highest predicted default risk and bin 5 to the 24.9% of all borrowers with the lowest predicted default risk. Finally, we calculate the fraction of borrowers in each credit score category that is in each of the 5 predicted default risk categories and their realized and predicted default rate. The results are displayed in Table 13. We also report the average realized and predicted default rate for each

<sup>20</sup>Appendix E also plots the time series of the rank correlation and the Gini coefficient for the credit score and our model for the current population. The credit score shows a large drop in these statistics for the credit score during the Great Recession whereas for our model they are both stable over time. This is consistent with the notion that the performance of the credit score dropped during the 2007-2009 period.

<sup>21</sup>The threshold levels for these categories are: 1) Deep Subprime: up to 499 credit score; 2) Subprime: 500-600 credit score; 3) Near Prime: 601-660 credit score; 4) Prime: 661-780 credit score; 5) Super Prime: higher than 781 credit score.



Table 12: Credit Score-Predicted Probability of Default Comparison– Current Borrowers

Metric	2006	2007	2008	2009	2010	2011	2012	2013
<b>Rank Correlation</b>								
Credit Score	0.9670	0.9613	0.9445	0.9653	0.9683	0.9585	0.9806	0.9559
Predicted Probability	0.9978	0.9979	0.9979	0.9975	0.9975	0.9973	0.9971	0.9974
Credit Score	0.6933	0.6935	0.6908	0.6807	0.6777	0.6833	0.6795	0.6810
Predicted Probability	0.7366	0.7244	0.7216	0.7187	0.7238	0.7335	0.7360	0.7391

Rank correlation between the credit score, predicted probability of default according to our model and subsequent realized default frequency by year for the population of current borrowers. Current borrowers do not have any delinquencies. For the credit score, report the rank correlation between each unique value of the score and the default frequency. For predicted probability of default based on our hybrid model, we first generate a number of bins equal to the number of unique credit scores in the data and then calculate the realized default frequency for each bin. Source: Authors' calculations based on Experian Data.

credit score category overall (columns 7 and 8) and for each predicted default risk category for all credit score (last 5 rows).

These results suggest that our model does well in predicting the default probability of borrowers in all categories, with a slight tendency to under-predict the probability of default by 1-2 percentage points for the Deep Subprime and Subprime borrowers. The majority of Deep Subprime borrowers fall in the two lowest categories of predicted default risk. For Subprime borrowers, 64% fall into the corresponding second category of default risk, while 15% fall in the first and 17% into the third. This corresponds to a sizable discrepancy as the average realized default probability for Subprime borrowers is 79%, whereas it is 99% for those in the first category and only 44% for those in the third. By contrast, the predicted default risk is very close to the realized default risk for Subprime borrowers in all categories of the predicted default risk distribution, with a discrepancy under 1 percentage point for predicted default risk category 1 and 2, and 5 percentage points for category 3 and 4. Near Prime borrowers also display a wide dispersion across predicted default risk categories with only 43% falling into the corresponding third category, 24% falling in category 2 (higher default risk) and 31% falling into category 4 (lower default risk). Again, the realized default rates vary substantially for Near Prime borrowers by predicted default risk category, from 77% in category 2, to 41% and 20% in category 3 and 4, respectively, while the predicted default risk is much closer to the realized, with a maximum 4 percentage point discrepancy. The discrepancy in classification for the credit score are lower for Prime and Super Prime borrowers. 13% of Prime borrowers fall into category 3 (higher default risk), 13% in category 5 (lower default risk) and 71% in the corresponding category 4. The realized default rates are 11% for Prime borrowers in category 4, and 34% and 3% respectively for Prime borrowers in category 3 and 5. Only 18% of Super Prime borrowers fall in category 4 of predicted default risk (higher risk) and 82% fall in the corresponding category 5. Moreover, the differences in realized default risk between these categories are minor,

with a realized default rate of 5% and 2% for categories 4 and 5, respectively. These results suggest that credit scores misclassify borrowers across risk categories with very different realized default rates. By contrast, as shown in the bottom 5 rows of Table 13 and by columns (6) and (8), our model is very successful at predicting the default rate for borrowers irrespective of their credit score.

### 6.1.1 Coverage

Consumers with limited credit histories encounter severe difficulties in accessing credit markets. As explained earlier in the paper, lenders often rely on credit scores to make lending decisions. If a borrower’s credit report does not have sufficient information to evaluate their default risk, lenders are unlikely to grant credit. Consequently, consumers with limited credit histories have a hard time accessing credit markets. These consumers can be divided into two groups. The first group consists of individuals without credit records, often referred to in the industry as “credit invisibles.” The second group are those consumers who do have credit records, but are considered “unscorable,” that is they have insufficient credit histories to generate a credit score. Brevoort, Grimm, and Kambara (2016) find that 11% of the US population lacks credit records and an additional 8.3% have a credit record but do not have a credit score.

A credit record may be considered unscorable for two reasons: (1) it contains insufficient information to generate a score, meaning the record either has too few accounts or has accounts that are too new to contain sufficient payment history to calculate a reliable credit score; or (2) it has become stale in that it contains no recently reported activity. The exact definition of what constitutes insufficient or stale information differs across credit scoring models, as each model uses its own proprietary definition.<sup>22</sup>

The challenges that credit invisibles and unscored consumers face in accessing credit markets has generated considerable attention from researchers and industry participants. Brevoort, Grimm, and Kambara (2016) show that young, minority and low income borrowers are disproportionately represented among the unscored. Several studies have explored the potential of various types of alternative data to supplement the information contained in credit reports and allow credit scores to be generated for these consumers.<sup>23</sup> Our model generates a predicted probability of default for every individual in our sample without an empty credit record, so effectively there are no active borrowers that we do not score. Our ability to score every active borrower is partly due to the fact that our model does not use any lagged observations. As previously discussed, many of the features in our model have a temporal dimension which renders the use of lags unnecessary. This

---

<sup>22</sup>The FICO score has the most restrictive requirements and does not score borrowers who show no updates or reports on credit file in past 6 months, or no accounts at least 6 months old. The rest of the industry has been trying to expand the universe of scorable borrowers and typically adopts a more flexible approach to increase the ranks of borrowers who are scored. Vantage Score has been particularly pushing the need to expand the universe of scorable customers and they have implemented several successful changes to the most recent version of their scoring model that have indeed substantially decreased the number of unscored consumers. For more information, see <https://www.vantagescore.com/resource/174/scoring-credit-invisibles-using-machine-learning-techniques>.

<sup>23</sup>See for example Jagtiani and Lemieux (2018). For an industry perspective, see Carroll and Rehmani (2017).

Table 13: Comparison with Credit Score

Credit Score (1)	Share (2)	Predicted Default Probability		Default Rate		Average Default Rate	
		(3)	Share (4)	Realized (5)	Predicted (6)	Realized (7)	Predicted (8)
Deep Subprime	6.48%	1	47.83%	99.55%	99.49%	95.45%	94.59%
		2	51.21%	92.24%	90.85%		
		3	0.90%	63.66%	52.25%		
		4	0.04%	43.72%	12.97%		
		5	0.01%	34.04%	1.95%		
Subprime	21.22%	1	14.90%	99.35%	99.35%	78.64%	77.27%
		2	64.41%	84.32%	83.85%		
		3	16.67%	51.93%	46.41%		
		4	4.00%	21.78%	18.07%		
		5	0.02%	16.18%	2.13%		
Near Prime	14.09%	1	1.34%	99.11%	99.21%	43.71%	42.61%
		2	24.19%	76.56%	78.58%		
		3	42.94%	41.39%	40.38%		
		4	30.90%	19.63%	15.91%		
		5	0.63%	3.13%	2.56%		
Prime	33.31%	1	0.08%	98.85%	99.18%	14.31%	14.37%
		2	2.42%	74.86%	78.37%		
		3	13.17%	33.54%	36.31%		
		4	70.92%	10.67%	10.27%		
		5	13.40%	3.18%	2.46%		
Super Prime	24.90%	1	0.00%	100.00%	99.31%	2.56%	2.66%
		2	0.09%	81.36%	81.77%		
		3	0.21%	32.13%	34.30%		
		4	18.00%	5.43%	5.95%		
		5	81.69%	1.76%	1.76%		
All	100	1	99.44%	99.41%			
		2	83.95%	83.88%			
		3	41.65%	40.65%			
		4	11.42%	10.62%			
		5	2.02%	1.89%			

Borrowers are classified into 5 categories of default risk standard in the industry named in column (1). The threshold levels for these categories are: 1) Deep Subprime: up to 499 credit score; 2) Subprime: 500-600 credit score; 3) Near Prime: 601-660 credit score; 4) Prime: 661-780 credit score; 5) Super Prime: higher than 781 credit score. The fraction of borrowers in each category is reported in column (2). Borrowers are also assigned to 5 categories of predicted default risk based on our hybrid model from the highest default risk (1) to the lowest (5), where the share of borrowers in each predicted default risk category is the same as for the credit score categories Deep Subprime to Super Prime. For each credit score risk category the share of borrowers in each predicted default risk category is reported in column (4). Columns (5) and (6) report the corresponding realized and predicted default probability for each credit score category interacted with predicted default risk category. Columns (7) and (8) report the average realized and predicted default probability for each credit score category. All rates, fractions and shares in percentage. Total # of observations: 17,732,772. Time period 2006Q1-2013Q4. Source: Authors' calculations based on Experian Data.

constitutes an additional advantage of our model when compared to traditional credit scores.

## 6.2 Predicting Systemic Risk

We next turn to analyze the aggregate forecasting power of our hybrid model. We aggregate the deep-learning forecasts for individual accounts to generate macroeconomic forecasts of credit risk by taking the average of the predicted probabilities over a given forecast period. Since our sample of consumers is nationally representative in each quarter, this will provide an unbiased estimate of the aggregate default risk predicted by our model. We calculate the aggregate default probability for 2006Q1-2013Q4, and show that our model is able to predict the spike in delinquencies during the 2007-2009 financial crisis and also the reduction in delinquencies since then. This estimate of aggregate default risk could be used as a proxy of systemic risk in the household sector. The results are displayed in Figure 11. Panel (a) plots the aggregate predicted default rate from our hybrid model and compares it to the aggregate realized default rate. While our predicted aggregate default rate is approximately 2 percentage points lower than the realized in 2006 and 2007, it rises at a similar speed as the realized default rate. It peaks in 2010Q2, approximately 2 quarters after the peak in the realized rate and then declines in the ensuing period, again reflecting the behavior of the realized rate, though it overestimates it by about 1 percentage point. Panel (b) shows a scatter plot of the predicted aggregate default rate against the realized for the different quarters in our sample period. The correlation between the predicted and realized aggregate default rate is 36%.

## 6.3 Value Added

We assess the economic salience of our hybrid DNN-GBT model by analyzing its value added for lenders and borrowers. For lenders, we examine the role our model can play in minimizing the losses from default. For borrowers, we calculate the interest savings for borrowers who are misclassified as having an excessively high probability of default based on the credit score compared to our model.

### 6.3.1 Lenders

We follow the framework proposed by Khandani, Kim, and Lo (2010), which compares the value of having a prediction of default risk to having none, and we make the same simplifying assumptions with respect to the revenues and costs of the consumer lending business. Specifically, in absence of any forecasts, it is assumed a lender will take no action regarding credit risk, implying that customers who default will generate losses for the lender, and customers who are current on their payments will generate positive revenues from financing fees on their running balances. To simplify, we assume that all defaulting and non-defaulting customers have the same running balance,  $B_r$ , but defaulting customers increase their balance to  $B_d$  prior to default. We refer to the ratio between  $B_d$  and  $B_r$  as "run-up." It is assumed that with a model to predict default risk, a lender can avoid losses of defaulting customers by cutting their credit line and avoiding run-up. Then, the value

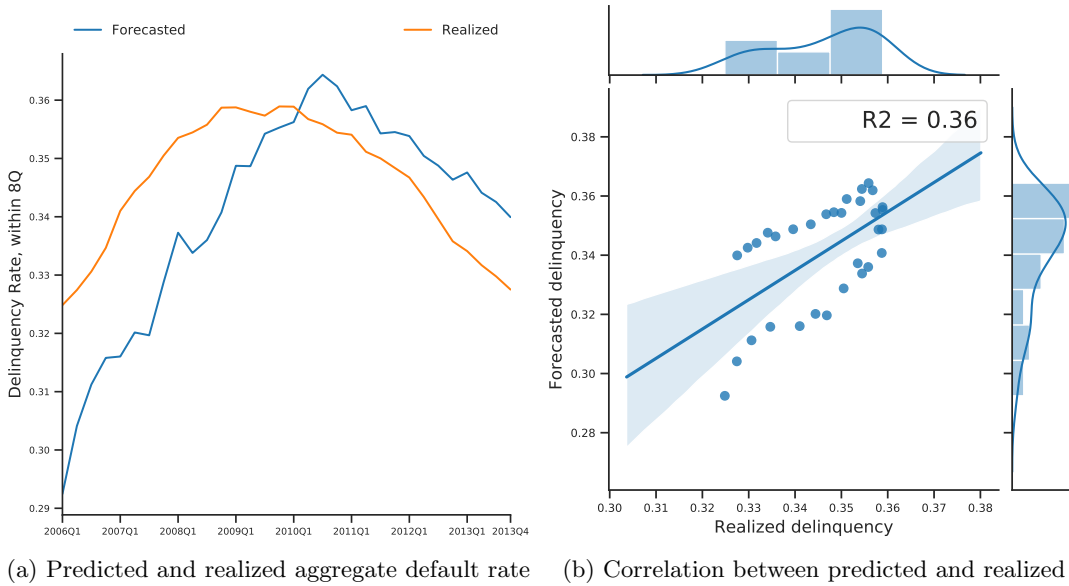


Figure 11: Fraction of consumers with 90+ days delinquency within the subsequent 8 quarters, predicted by our hybrid DNN-GBT model and realized. Aggregate default rates are obtained by averaging across all consumers in each period. Source: Authors' calculations based on Experian Data.

added as proposed by Khandani, Kim, and Lo (2010) can be written as follows:

$$VA(r, N, TN, FN, FP) = \frac{TN - FN[1 - (1 + r)^{-N}] \left[ \frac{B_d}{B_r} - 1 \right]^{-1}}{TN + FP} \quad (10)$$

where  $r$  refers to the interest rate,  $N$  the loan's amortization period, and  $TN, FN, FP$  refer to true negatives, false negatives and false positives respectively. Panel (a) of Figure 12 plots the Value Added (VA) as a function of interest rate and the ratio of run-up balance for our out-of-sample forecasts of 90+ days delinquencies over the subsequent 8 quarters for 2012Q4. These estimates imply cost savings of over 60% of total losses when compared to having no forecast model for a run-up of 1.2 at a 10% interest rate for an amortization period for 3 years.

We next compare the value added of our hybrid model with default predictions generated by a logistic regression. This exercise illustrates the gains from adopting a better technology for credit allocation. Panel (b) of Figure 12 shows more modest, but substantial cost savings in the range of 1-9% and approximately 5% for a 1.2 run-up at a 10% interest rate with a 3 year amortization period. This exercise then confirms the advantages of using deep learning over other technologies in predicting default.




Figure 12: Value-added of machine-learning forecasts of 90+ days delinquency over 8Q forecast horizons on data from 2012Q4. VA values are calculated with amortization period  $N = 3$  years and a 50% classification threshold. Source: Authors' calculations based on Experian Data.

### 6.3.2 Borrowers

We now examine the potential cost savings for consumers who would be offered credit according to the predicted default probability implied by our model instead of a conventional credit score. Following our approach in Section 6.1, we create credit score categories based on common industry standards and corresponding predicted probability bins with the same number of observations, and we place customers in these bins based on their credit score at account origination for each of their credit cards.<sup>24</sup> The distribution of customers is summarized in Table 24 in Appendix F. We then follow the information on interest rates by credit score category in Table 2 in Agarwal et al. (2015) to obtain the cost of credit on credit card balances.<sup>25</sup> To obtain the cost savings for consumers, we use the difference in interest rates by credit score category based on how they would be classified according to our model. For customers who are placed in higher risk categories by the credit score compared to our predicted probability of default, interest rates on credit cards are higher than they would have been if they had been classified according to our model. Thus, using our model to score consumers rather than the credit score would generate the cost savings for them. For customers placed in risk categories by the credit score that are too low relative to the default risk predicted by our model, interest rates will be higher under our model. The calculation is made for each individual consumer. The average for each credit score category is then computed and reported in Table 25. The information on interest rates and balances, and the dollar value of cost savings for

<sup>24</sup>We look at the months since most recently opened credit card account to infer account origination. We drop customers with months since the most recently opened credit card greater than 72.

<sup>25</sup>Credit card interest rates are notoriously invariant to overall changes in interest rates, so the calculations reported in this section apply irrespective of the time period. See Ausubel (1991) and Calem and Mester (1995).

different credit card categories is reported in Table 14. We report this in percentage of credit card balances in the top panel and in current USD terms in the bottom panel. The largest gains accrue to customers with Subprime and Near Prime credit scores. As we showed in Section 6.1, they are more likely to be attributed a probability of default by the credit score that is too low compared to our model predictions. Additionally, the biggest variation in credit card interest rates occurs across Subprime and Near Prime borrowers in comparison to Prime based on Agarwal et al. (2015). The cost savings for these borrowers average out to 4-5% of total credit card balances of \$1,085-1426. Gains for Prime and Super Prime borrowers who are attributed a lower default probability by our model are very modest, as credit card interest rates vary little by credit score for Prime and Superprime borrowers. On the other hand, Prime and Superprime borrowers who are placed in group 1, corresponding to the highest predicted default probability based on our model face, substantial losses in the order of 4-5% of total credit card balances or \$274-423. The cumulated interest rate cost savings across all consumers in our sample is \$723,636,560, which amounts to \$40 per capita.

This calculation provide us with a lower bound for the cost savings of being classified according to our model in comparison to the credit score, as they do not take into account the higher credit limits and potential behavioral responses of customers faced with higher borrowing capacity and lower interest rates. As shown in Agarwal et al. (2015), changes in the cost of funds for lenders mainly translate into changes in credit limits and exclusively for higher credit score borrowers. Therefore, being placed in a higher risk category for consumers also inhibits their ability to benefit from expansionary monetary policy. Additionally, we do not take into account the fact that more expensive credit in the form of higher interest rate costs makes it more likely that the consumer will incur missed payments in response to temporary changes in income. Fees for missed payments constitute a substantial component of credit card costs for consumers, and the ability to avoid these fees would contribute to substantial cost savings for consumers (see Agarwal et al. (2014)).

## 7 Conclusion

We have proposed to use deep learning to develop a model to predict consumer default. Our model uses the same data used by conventional scoring models and abides with all legislative restrictions in the United States. We show that our model compares favorably to conventional credit scoring models in ranking individual consumers by their default risk, and is also able to capture variations in aggregate default risk. Our model is interpretable and allows to identify the factors that are most strongly associated with default. Whereas conventional credit scoring models emphasize utilization rates, our analysis suggests that the number and balances on open trades are the factors which associate more strongly to higher default probabilities. Our model is able to provide a default prediction for all consumers with a non-empty credit record. Additionally, we show that our hybrid DNN-GBT model performs better than standard machine learning models of default based on

Table 14: Cost of Credit Risk Misclassification

		Credit Score			
		Subprime, Near Prime	Prime Low	Prime Mid	Prime High, Superprime
<b>Annual Interest Rate Savings</b>					
Predicted Default Bin	1	0.00%	-5.13%	-4.28%	-4.93%
	2	5.13%	0.00%	0.85%	0.20%
	3	4.28%	-0.85%	0.00%	-0.65%
	4	4.93%	-0.20%	0.65%	0.00%
<b>Annual Average Cost Saving (\$)</b>					
Predicted Default Bin	1	0	-423	-274	-309
	2	1085	0	83	15
	3	1426	-179	0	-54
	4	1239	-44	118	0

This table reports the average cost savings for consumers across credit score and predicted default probability bins for our sample. The cumulative savings for consumers on both credit card and bankcard debt adds up to \$723,636,560. Time period 2006Q1-2013Q4. Source: Authors' calculations based on Experian Data.

logistic regression and can accrue cost saving to lenders in the order of 1-9% compared to default predictions based on logistic regression, as well as interest rate cost savings for consumers of up to \$1,426 per year.



## References

- Abadi, Martín, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. “Tensorflow: a system for large-scale machine learning.” *OSDI*, Volume 16. 265–283.
- Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel. 2014. “Regulating consumer financial products: Evidence from credit cards.” *The Quarterly Journal of Economics* 130 (1): 111–164.
- . 2015, September. “Do Banks Pass Through Credit Expansions to Consumers Who Want to Borrow?” Working paper 21567, National Bureau of Economic Research.
- Albanesi, Stefania, Giacomo De Giorgi, and Jaromir Nosal. 2017, August. “Credit Growth and the Financial Crisis: A New Narrative.” Working paper 23740, National Bureau of Economic Research.
- Athey, Susan, and Guido W Imbens. 2019. “Machine learning methods that economists should know about.” *Annual Review of Economics*, vol. 11.
- Athreya, Kartik, Xuan S Tam, and Eric R Young. 2012. “A quantitative theory of information and unsecured credit.” *American Economic Journal: Macroeconomics* 4 (3): 153–83.
- Ausubel, Lawrence M. 1991. “The failure of competition in the credit card market.” *The American Economic Review*, pp. 50–81.
- Bengio, Yoshua, and Yann LeCun. 2007. “Scaling Learning Algorithms Towards AI.” In *Large Scale Kernel Machines*, edited by Léon Bottou, Olivier Chapelle, D. DeCoste, and J. Weston. MIT Press.
- Bergstra, James S, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. 2011. “Algorithms for hyperparameter optimization.” *Advances in neural information processing systems*. 2546–2554.
- Breiman, Leo. 2001. “Random forests.” *Machine learning* 45 (1): 5–32.
- Brevoort, Kenneth P., Philipp Grimm, and Michelle Kambara. 2016. “Credit Invisibles and the Unscored.” *Cityscape* 18 (2): 9–34.
- Butaru, Florentin, Qingqing Chen, Brian Clark, Sanmay Das, Andrew W Lo, and Akhtar Siddique. 2016. “Risk and risk management in the credit card industry.” *Journal of Banking & Finance* 72:218–239.
- Calem, Paul S, and Loretta J Mester. 1995. “Consumer behavior and the stickiness of credit-card interest rates.” *The American Economic Review* 85 (5): 1327–1336.
- Carroll, Peter, and Saba Rehmani. 2017. “Alternative Data and the Unbanked.” *Oliver Wyman Insights*.

- Chatterjee, Satyajit, Dean Corbae, Kyle Dempsey, and José-Victor Ríos-Rull. 2016. “A theory of credit scoring and competitive pricing of default risk.” *Unpublished paper, University of Wisconsin*, vol. 31.
- Chatterjee, Satyajit, Dean Corbae, Makoto Nakajima, and José-Víctor Ríos-Rull. 2007. “A quantitative theory of unsecured consumer credit with risk of default.” *Econometrica* 75 (6): 1525–1589.
- Chen, Hugh, Scott Lundberg, and Su-In Lee. 2018. “Hybrid Gradient Boosting Trees and Neural Networks for Forecasting Operating Room Data.” *arXiv preprint arXiv:1801.07384*.
- Chen, Tianqi, and Carlos Guestrin. 2016. “XGBoost: A Scalable Tree Boosting System.” *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16. New York, NY, USA: ACM, 785–794.
- Chollet, François, et al. 2015. “Keras: Deep learning library for theano and tensorflow.” *URL: https://keras.io/k* 7, no. 8.
- Corbae, Dean, and Andrew Glover. 2018, September. “Employer Credit Checks: Poverty Traps versus Matching Efficiency.” Working paper 25005, National Bureau of Economic Research.
- Friedman, Jerome H. 2001. “Greedy function approximation: a gradient boosting machine.” *Annals of statistics*, pp. 1189–1232.
- Goodfellow, Ian, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. 2016. *Deep learning*. Volume 1. MIT press Cambridge.
- Hornik, Kurt. 1991. “Approximation capabilities of multilayer feedforward networks.” *Neural Networks* 4 (2): 251 – 257.
- Hornik, Kurt, Maxwell Stinchcombe, and Halbert White. 1989. “Multilayer feedforward networks are universal approximators.” *Neural Networks* 2 (5): 359 – 366.
- Hunter, John D. 2007. “Matplotlib: A 2D graphics environment.” *Computing in science & engineering* 9 (3): 90–95.
- Ioffe, Sergey, and Christian Szegedy. 2015. “Batch normalization: Accelerating deep network training by reducing internal covariate shift.” *arXiv preprint arXiv:1502.03167*.
- Jagtiani, Julapa, and Catharine Lemieux. 2018. “The roles of alternative data and machine learning in fintech lending: evidence from the LendingClub consumer platform.”
- Khandani, Amir E, Adlar J Kim, and Andrew W Lo. 2010. “Consumer credit-risk models via machine-learning algorithms.” *Journal of Banking & Finance* 34 (11): 2767–2787.
- Kingma, Diederik P, and Jimmy Ba. 2014. “Adam: A method for stochastic optimization.” *arXiv preprint arXiv:1412.6980*.

- Kuncheva, Ludmila I, and Christopher J Whitaker. 2003. “Measures of diversity in classifier ensembles and their relationship with the ensemble accuracy.” *Machine learning* 51 (2): 181–207.
- Kvamme, Håvard, Nikolai Sellereite, Kjersti Aas, and Steffen Sjursen. 2018. “Predicting mortgage default using convolutional neural networks.” *Expert Systems with Applications* 102:207–217.
- Lessmann, Stefan, Bart Baesens, Hsin-Vonn Seow, and Lyn C Thomas. 2015. “Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research.” *European Journal of Operational Research* 247 (1): 124–136.
- Livshits, Igor, James MacGee, and Michele Tertilt. 2007. “Consumer bankruptcy: A fresh start.” *American Economic Review* 97 (1): 402–418.
- Lundberg, Scott M, and Su-In Lee. 2017. “A Unified Approach to Interpreting Model Predictions.” In *Advances in Neural Information Processing Systems 30*, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, 4765–4774. Curran Associates, Inc.
- McKinney, Wes, et al. 2010. “Data structures for statistical computing in python.” *Proceedings of the 9th Python in Science Conference*, Volume 445. Austin, TX, 51–56.
- Molnar, Christoph. 2019. *Interpretable Machine Learning*. <https://christophm.github.io/interpretable-ml-book/>.
- Montúfar, Guido, Razvan Pascanu, Kyunghyun Cho, and Yoshua Bengio. 2014. “On the Number of Linear Regions of Deep Neural Networks.” *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, NIPS’14. Cambridge, MA, USA: MIT Press, 2924–2932.
- Mullainathan, Sendhil, and Jann Spiess. 2017. “Machine learning: an applied econometric approach.” *Journal of Economic Perspectives* 31 (2): 87–106.
- Pedregosa, Fabian, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. “Scikit-learn: Machine learning in Python.” *Journal of machine learning research* 12 (Oct): 2825–2830.
- Ren, Xudie, Haonan Guo, Shenghong Li, Shilin Wang, and Jianhua Li. 2017. “A novel image classification method with CNN-XGBoost model.” *International Workshop on Digital Watermarking*. Springer, 378–390.
- Schmidhuber, Jürgen. 2015. “Deep learning in neural networks: An overview.” *Neural networks* 61:85–117.
- Sirignano, Justin, Apaar Sadhwani, and Kay Giesecke. 2018. “Deep Learning for Mortgage Risk.” *Available at SSRN*.

- Srivastava, Nitish, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. “Dropout: a simple way to prevent neural networks from overfitting.” *The Journal of Machine Learning Research* 15 (1): 1929–1958.
- Walt, Stéfan van der, S Chris Colbert, and Gael Varoquaux. 2011. “The NumPy array: a structure for efficient numerical computation.” *Computing in Science & Engineering* 13 (2): 22–30.
- West, David. 2000. “Neural network credit scoring models.” *Computers & Operations Research* 27 (11-12): 1131–1152.
- Xia, Yufei, Chuanzhe Liu, YuYing Li, and Nana Liu. 2017. “A boosted decision tree approach using Bayesian hyper-parameter optimization for credit scoring.” *Expert Systems with Applications* 78:225–241.

# Appendix

## A Performance Metrics

Suppose a binary classifier is given and applied to a sample of  $N$  observations. For each instance  $i$ , let  $y_i$  denote the true outcome. For each observation, the model generates a probability that an observation with feature vector  $x_i$  belongs to class 1. This predicted probability,  $f(x_i)$  is then evaluated based on a threshold to classify observations into class 1 or 0. Given a threshold level ( $c$ ), let True Positive (TP) denote the number of observations that are correctly classified as type 0, True Negative (TN) be the number of observations that are correctly classified as type 1, False Positive (FP) be the number of observations that are type 1 but incorrectly classified as type 0, and, finally, False Negative (FN) be the number of observations that are actually of type 0 but incorrectly classified as type 0. Based on these definitions, one can define the following metrics to assess the performance of the classifier:

$$\text{True Positive Rate (TPR)} \equiv \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (11)$$

$$\text{False Positive Rate (FPR)} \equiv \frac{\text{FP}}{\text{FP} + \text{TN}} \quad (12)$$

$$\text{Precision} \equiv \frac{\text{TN}}{\text{TN} + \text{FN}} \quad (13)$$

$$\text{Recall} \equiv \frac{\text{TN}}{\text{TN} + \text{FP}} \quad (14)$$

$$\text{F-measure} \equiv \frac{2 \times \text{Recall} \times \text{Precision}}{\text{Precision} + \text{Recall}} \quad (15)$$

$$\text{Accuracy} \equiv \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \quad (16)$$

$$\text{Youden's J statistic} \equiv \frac{\text{TP}}{\text{TP} + \text{FN}} + \frac{\text{TN}}{\text{TN} + \text{FP}} - 1 \quad (17)$$

$$\text{ROC AUC} = \int_{-\infty}^{\infty} \text{TPR}(c) \text{FPR}'(c) dc \quad (18)$$

$$\text{Cross-entropy loss} = -\frac{1}{N} \sum_{i=1}^N (y_i \cdot \log(f(x_i)) + (1 - y_i) \cdot \log(1 - f(x_i))) \quad (19)$$

$$\text{Brier score} = \frac{1}{N} \sum_{i=1}^N (f(x_i) - y_i)^2 \quad (20)$$

## B Data Pre-Processing

Our original dataset contains 33,600,000 observations. We discard observations of individuals with missing birth information, deceased individuals and restrict our analysis to individuals aged between 18 and 85, residing in one of the 50 states or the District of Columbia, with 8 consecutive quarters of non-missing default behavior. This leaves us with 22,004,753 data points. Our itemized sample

restrictions are summarized in Table 15 below.

Table 15: Itemized Sample Restrictions

	Observations
Credit Report Data	33,600,000
Remove	
Deceased	- 513,270
Age	- 4,718,804
Residence	- 953,215
Prediction Window	- 5,409,958
Prediction Sample	22,004,753

## Feature Scaling

We normalize all explanatory variables by their means and standard deviations:

$$z_i = \frac{x_i - \mu_x}{\sigma_x} \tag{21}$$

where  $\mathbf{x} = (x_1, x_2, \dots, x_k)$ , and  $\mathbf{z}_i$  is the  $i^{\text{th}}$  normalized data.

## Train-Test Split

For most of our analysis we split the data to account for look-ahead bias, i.e., the training set consists of data 8Q prior to the testing data. Then, we scale the testing data by the mean and standard deviation of the training data. In an alternative specification, we split our pooled data into three chunks: training set (60%), holdout set (20%), and testing set (20%). We report each specifications in Table 3 - Table 4. Except for parts of Section 5.2, we used the predictions generated by our models on the temporal splits.

In each specifications, we randomly shuffled the data to ensure that the mini-batch gradients are unbiased. If gradients are biased, training may not converge and accuracy may be lost.

## C Model Estimation

Our estimation consists of four steps. First, we specify the loss function. Second, we choose the optimization algorithm. Third, we optimize the hyperparameters of the model. Fourth, we train our models.

## Loss Function

Suppose  $\mathbf{y}$  is the ground truth vector of default, and  $\hat{\mathbf{y}}$  is the estimate obtained directly from the last layer given input vector  $\mathbf{x} = (x_1, x_2, \dots, x_k)$ . By construction,  $y_i = \{0, 1\}$  and  $\hat{y}_i \in [0, 1]$ . We minimize the categorical cross-entropy loss function<sup>26</sup> to estimate the parameter specified in (7). We do this by choosing  $\theta$  that minimizes the distance between the predicted  $\hat{\mathbf{y}}$  and the actual  $\mathbf{y}$  values. Given  $N$  training examples, the categorical cross-entropy loss can be written as:

$$L(\hat{y}, y) = -\frac{1}{N} \sum_{i=1}^N (y_i \cdot \log(\hat{y}_i) + (1 - y_i) \cdot \log(1 - \hat{y}_i)) \quad (22)$$

We apply an iterative optimization algorithm to find the minimum of the categorical cross-entropy loss function. We next describe this algorithm.

## DNN Optimization Algorithm

Deep learning models are computationally demanding due to their high degree of non-linearity, non-convexity and rich parameterization. Given the size of the data, gradient descent is impractical. We follow the standard approach of using stochastic gradient descent (SGD) to train our deep learning models (see Goodfellow et al. (2016)). Stochastic gradient descent is an iterative algorithm that uses small random subsets of the data to calculate the gradient of the objective function. Specifically, a subset of the data, referred to as a mini-batch (the size of the mini-batch is called the batch size), is loaded into memory and the gradient is computed on this subset. The gradient is then updated, and the process is repeated until convergence.

We adopt the Adaptive Moment Estimation (Adam), a computationally efficient variant of the SGD introduced by (see Kingma and Ba (2014)) to train our neural networks. The Adam optimization algorithm can be summarized as follows:

1. Fix the learning rate  $\alpha$ , the exponential decay rates for the moment estimates:  $\beta_1, \beta_2 \in [0, 1)$ , and the objective function. Initialize the parameter vector  $\theta_0$ , the first and second moment vector  $m_0$  and  $v_0$  respectively, and the timestep  $t$ .
2. While  $\theta_t$  does not converge, do the following:
  - (a) Compute the gradients with respect to the objective function at timestep  $t$ :

$$g_t = \nabla_{\theta} f_t(\theta_{t-1}) \quad (23)$$

- (b) Update the first and second moment estimates:

$$m_t = \beta_1 \cdot m_{t-1} + (1 - \beta_1) \cdot g_t \quad (24)$$

$$v_t = \beta_2 \cdot v_{t-1} + (1 - \beta_2) \cdot g_t^2 \quad (25)$$

---

<sup>26</sup>Loss function measures the inconsistency between the predicted and the actual value. The performance of a model increases as the loss function decreases. There are several other types of loss functions, including mean squared error, hinge, and Poisson. The categorical cross-entropy is often used for classification problems.

(c) Compute the bias-corrected first and second moment estimates:

$$\hat{m}_t = \frac{m_t}{1 - \beta_1^t} \quad (26)$$

$$\hat{v}_t = \frac{v_t}{1 - \beta_2^t} \quad (27)$$

(d) Update the parameters:

$$\theta_t = \theta_{t-1} - \frac{\alpha \cdot \hat{m}_t}{\sqrt{\hat{v}_t} + \epsilon} \quad (28)$$

The hyperparameters have intuitive interpretations and typically require little tuning. We apply the default setting suggested by the authors of Kingma and Ba (2014), these are  $\alpha = 0.001$ ,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$  and  $\epsilon = 10^{-7}$ .

## GBT Algorithm

Fit a shallow tree (e.g., with depth  $L = 1$ ). Using the prediction residuals from the first tree, fit a second tree with the same shallow depth  $L$ . Weight the predictions of the second tree by  $\nu \in (0, 1)$  to prevent the model from overfitting the residuals, and then aggregate the forecasts of these two trees. Until a total of  $K$  trees is reached in the ensemble, at each step  $k$ , fit a shallow tree to the residuals from the model with  $k-1$  trees, and add its prediction to the forecast of the ensemble with a shrinkage weight of  $\nu$ .

## Regularization

Neural networks are low-bias, high-variance models (i.e., they tend to overfit to their training data). We implement three routines to mitigate this. First, we apply dropout to each of the layers (see Srivastava et al. (2014)). During training, neurons are randomly dropped (along with their connections) from the neural network with probability  $p$  (referred to as the dropout rate), which prevents complex co-adaptations on training data.

Second, we implement "early stopping", a general machine learning regularization tool. After each time the optimization algorithm passes through the training data (i.e., referred to as an epoch), the parameters are gradually updated to minimize the prediction errors in the training data, and predictions are generated for the validation sample. We terminate the optimization when the validation sample loss has not decreased in the past 50 epochs. Early stopping is a popular substitute to l2 regularization, since it achieves regularization at a substantially lower computational cost.

Last, we use batch normalization (see Ioffe and Szegedy (2015)), a technique for controlling the variability of features across different regions of the network and across different datasets. It is motivated by the internal covariate shift, a phenomenon in which inputs of hidden layers follow different distributions than their counterparts in the validation sample. This problem is frequently encountered when fitting deep neural networks that involve many parameters and rather complex structures. For each hidden unit in each training step, the algorithm cross-sectionally de-means and variance standardizes the batch inputs to restore the representation power of the unit.



## Hyperparameter selection

Deep learning models require a number of hyperparameters to be selected. We follow the standard approach by cross-validating the hyperparameters via a validation set. We fix a training and validation set, and then train neural networks with different hyperparameters on the training set and compare the loss function on the validation set. We cross-validate the number of layers, the number of units per layer, the dropout rate, the batch size, and the activation function (i.e., the type of non-linearity) via Tree-structured Parzen Estimator (TPE) approach (see Bergstra et al. (2011)),<sup>27</sup> and select the hyperparameters with the lowest validation loss.

The training set for our out-of-sample hyperparameter optimization comes from 2004Q3, while the validation set is from 2006Q3. Table 16 summarizes our machine learning model hyperparameters. For our neural network, we used 5 hidden layers, with 150-600-1000-600-400 neurons per layer, SELU activation function, a batch size of 4096, a learning rate of 0.003, and a dropout rate of 50%. For our GBT, We found that a learning rate of 0.05, a max tree depth of 6, a max bin size of 64, with 1000 trees gave us good performance. All GBT models were run until their validation accuracy was non-improving for a hundred rounds and were trained on CPUs.

Table 16: Hyperparameters for Machine Learning Models: Out-of-sample Exercise

Model	Tree Depth	# of Trees
CART	7	
RF	20	900
GBT	6	1000

For the pooled sample prediction, we increased the number of neurons per layers to 512,1024,2048,1024,512 and decreased the dropout rate to 20%, keeping the activation function, the batch size, and the learning rate unchanged. We instituted early stopping with a patience of 1,000 for GBT, and trained a model of depth 6 with up to 10,000 trees and a learning rate of 0.3. We report the results of the best performing GBT.

## Implementation

We include 139 features for each individual. Since we work with panel data, there is a sample for each quarter of data. We train roughly 20 million samples, which takes up around 20 gigabytes of data. Our deep learning models are made up of millions of free parameters. Since the estimation procedure relies on computing gradients via backpropagation, which tends to be time and memory intensive, using conventional computing resources (e.g., desktop) would be impractical (if not infeasible). We address the time and memory intensity with two methods. First, to save memory, we use single precision floating point operations, which halves the memory requirements and results in a substantial computational speedup. Second, to accelerate the learning, we parallelized our computations and trained all of our models on a GPU cluster<sup>28</sup>. In our setting, GPU computations

<sup>27</sup>We use TPE since it outperformed random search (see Bergstra et al. (2011)), which was shown to be both theoretically and empirically more efficient than standard techniques such as trials on a grid. Other widely used strategies are grid search and manual search.

<sup>28</sup>1 node with 4 NVIDIA GeForceGTX1080 GPUs. The pooled model trains within 36 hours.

were over 40X faster than CPU for our deep neural networks. For a discussion on the impact of GPUs in deep learning see Schmidhuber (2015).

We conduct our analysis using Python 3.6.3 (Python Software Foundation), building on the packages numpy (Walt, Colbert, and Varoquaux (2011)), pandas (McKinney et al. (2010)) and matplotlib (Hunter (2007)). We develop our deep neural networks with keras (Chollet et al. (2015)) running on top of Google TensorFlow, a powerful library for large-scale machine learning on heterogeneous systems (Abadi et al. (2016)). We run our machine learning algorithms using sci-kit learn (Pedregosa et al. (2011)) and (Chen and Guestrin (2016)).

## Features

Table 17 lists our model inputs. Table 18 provides summary statistics for selected features. For the SHAP value analysis, we grouped features that had a correlation higher than 0.7. These groups are presented in Table 19.

Table 17: Model Inputs

---

90 day delinquencies in the last 36 months	Credit amount paid down on open first mortgage trades
90 days delinquencies in the last 12 months	Credit amount paid down on open second mortgage trades
90 days delinquencies in the last 24 months	Credit card trades opened in the last 12 months
90 days delinquencies in the last 6 months	Credit card utilization ratio
Auto loan or lease inquiries made in the last 3 months	Dismissed bankruptcies
Auto loan trades opened in the last 6 months	Early payoff trades
Balance on 30 days late bankcard trades	Fannie Mae first mortgage trades opened prior to June 2009
Balance on 30 days late mortgage trades	First mortgage trades opened in the last 6 months
Balance on 60 days late bankcard trades	Fraction of 30 days delinquent debt to total debt
Balance on 60 days late mortgage trades	Fraction of 60 days delinquent debt to total debt
Balance on 90-180 days late bankcard trades	Fraction of 90 days delinquent debt to total debt
Balance on 90-180 days late mortgage trades	Fraction of auto loan to total debt
Balance on authorized user trades	Fraction of credit card debt to total debt
Balance on bankcard trades	Fraction of HELOC to total debt
Balance on collections	Fraction of mortgage to total debt
Balance on credit card trades	Freddie Mac first mortgage trades opened prior to June 2009
Balance on derogatory bankcard trades	HELOC trades ever 90 or more days delinquent or derogatory
Balance on derogatory mortgage trades	HELOC utilization ratio
Balance on first mortgage trades	Inquiries made in the last 12 months
Balance on HELOC trades	Installment trades
Balance on mortgage trades	Installment utilization ratio
Balance on open 30 days late installment trades	Judgments with amount >\$1000
Balance on open 30 days late revolving trades	Monthly payment on all debt
Balance on open 60 days late installment trades	Monthly payment on credit card trades
Balance on open 60 days late revolving trades	Monthly payment on HELOC trades
Balance on open 90-180 days late installment trades	Monthly payment on open auto loan trades
Balance on open 90-180 days late revolving trades	Monthly payment on open first mortgage trades
Balance on open auto loan trades	Monthly payment on open second mortgage trades
Balance on open bankcard trades with credit line suspended	Monthly payment on student trades
Balance on open derogatory installment trades	Mortgage trades
Balance on open derogatory revolving trades	Mortgage type
Balance on open installment trades	Mortgage inquiries made in the last 3 months
Balance on open personal liable business loans	Open auto loan trades
Balance on open revolving trades	Open bankcard trades
Balance on revolving trades	Open bankcard trades opened in the last 6 months
Balance on second mortgage trades	Open credit card trades
Balance on student trades	Open first mortgage trades
Bankcard inquiries made in the last 3 months	Open HELOC trades
Bankruptcies filed within the last 24 months	Open mortgage trades
Chapter 13 bankruptcies	Open personal liable business loans
Chapter 7 bankruptcies	Open second mortgage trades
Charge-off amount on unsatisfied charge-off trades	Petitioned bankruptcies
Charge-off trades	Public record bankruptcies
Collections placed in the last 12 months	Public records filed in the last 24 months
Credit amount on deferred student trades	Total 30 days late debt balances
Credit amount on non-deferred student trades	Total 60 days late debt balances
Credit amount on open credit card trades	Total 90 or more days delinquent debt balances
Credit amount on open HELOC trades	Total 90-180 days late debt balances
Credit amount on open installment trades	Total credit amount on open trades
Credit amount on open mortgage trades	Total debt balances
Credit amount on revolving trades	Total derogatory debt balances
Credit amount on unsatisfied derogatory trades	Trades legally paid in full for less than the full balance
:	:

---

---

:  
 Months since the most recently closed transferred or refinanced first mortgage  
 Months since the most recently opened auto loan trade  
 Months since the most recently opened credit card trade  
 Months since the most recently opened first mortgage  
 Months since the most recently opened HELOC trade  
 Months since the most recently opened second mortgage  
 Months since the most recent 30-180 days delinquency on auto loan or lease trades  
 Months since the most recent 30-180 days delinquency on mortgage trade  
 Months since the most recent 30-180days delinquency on credit card trades  
 Months since the most recent foreclosure proceeding started on first mortgage  
 Months since the most recent public record bankruptcy filed  
 Months since the most recent 30-180 days delinquency  
 Months since the oldest trade was opened  
 Months since the most recent 90 or more days delinquency  
 Presence of outstanding governmental agency debts  
 Presently foreclosed first mortgage trades that occurred in the last 24 months  
 Student trades ever 90 or more days delinquent or derogatory occurred in the last 24 months  
 Trades ever 90 or more days delinquent or derogatory occurred in the last 24 months  
 Presently foreclosed first mortgages  
 Ratio of inquiries to trades opened in the last 6 months  
 Utility trades  
 Utilization ratio  
 Unsatisfied collections  
 Unsatisfied repossession trades  
 Worst ever status on a credit card trade in the last 24 months  
 Worst ever status on a mortgage trade in the last 24 months  
 Worst ever status on an auto loan or lease trade in the last 24 months  
 Worst ever status on any trades in the last 24 months  
 Worst present status on a mortgage trade  
 Worst present status on a revolving trade  
 Worst present status on an auto loan or lease trade  
 Worst present status on an installment trade  
 Worst present status on an open trade  
 Worst present status on any trades  
 Worst present status on bankcard trades

---

List of features included in our model.

Table 18: Summary Statistics

Feature	Mean	Std. Dev	25%	Median	75%
Balance on collections	724.74	3951.88	0	0	0
Balance on credit card trades	4573.20	9824.31	0	834	4591
Balance on mortgage trades	63408.03	160889.35	0	0	76301
Balance on open auto loan trades	4472.88	11608.48	0	0	3915
Balance on open installment trades	8754.12	32554.07	0	0	10583
Balance on open personal liable business loan	290.94	17183.32	0	0	0
Balance on revolving trades	4532.12	9804.11	0	761	4478
Balance on student trades	3523.54	15537.59	0	0	0
Charge-off amount on unsatisfied charge-off trades	1264.56	83747.09	0	0	0
Collections placed in the last 12 months	0.38	1.24	0	0	0
Credit amount on open credit card trades	21475.90	30662.29	0	8600	31947
Credit amount on open installment trades	12178.34	37332.78	0	0	17286
Credit amount on revolving trades	21382.80	30396.75	0	8641	31890
Credit amount on unsatisfied derogatory trades	11048.38	70187.96	0	0	1500
Credit amount paid down on open first mortgage trades	6109.98	164527.99	0	0	2255
Credit card utilization ratio	0.51	2.31	0	0.06	0.37
Early payoff trades	0.94	1.64	0	0	1
Fraction of 30 days delinquent debt to total debt	0.02	0.11	0	0	0
Fraction of 60 days delinquent debt to total debt	0.01	0.08	0	0	0
Fraction of 90 days delinquent debt to total debt	0.03	0.15	0	0	0
Fraction of auto loan to total debt	0.12	0.27	0.00	0.00	0.04
Fraction of credit card debt to total debt	0.27	0.40	0.00	0.03	0.44
Fraction of home equity line of credit to total debt	0.03	0.13	0.00	0.00	0.00
Fraction of mortgage to total debt	0.30	0.42	0.00	0.00	0.82
Inquiries made in the last 12 months	1.38	2.28	0	1	2
Judgments with amount >\$1000	0.06	0.32	0	0	0
Monthly payment on all debt	907.95	11059.39	20	333	1232
Monthly payment on credit card trades	121.22	366.75	0	31	125
Monthly payment on open auto loan trades	132.70	313.24	0	0	227
Monthly payment on student trades	21.12	4628.97	0	0	0
Months since the most recently opened credit card trade	35.64	46.17	7	20	47
Months since the oldest trade was opened	196.45	126.35	98	178	271
Open auto loan trades	0.34	0.60	0	0	1
Open credit card trades	3.56	3.91	0	2	5
Open mortgage trades	0.50	0.83	0	0	1
Total 60 days late debt balances	619.09	13253.67	0	0	0
Total 90 or more days delinquent debt balances	3125.34	33186.26	0	0	0
Total credit amount on open trades	108480.31	259094.00	3000	33146	146535
Total debt balances	77126.36	170742.97	318	11738	95808
Total derogatory debt balances	1287.84	18422.40	0	0	0
Trades ever 90 or more days delinquent or derogatory, last 24 months	1.22	2.78	0	0	1
Utilization ratio	0.56	0.88	0.027	0.52	0.83

Table 19: Feature Groups

<u>Total debt balances*</u>	<u>Number of collections*</u>
Total debt balances	Collections placed in the last 12 months
Total credit amount on open trades	Trades ever 90 or more days delinquent
Balance on mortgage trades	Unsatisfied collections
Credit amount on open mortgage trades	
Balance on first mortgage trades	<u>Number of open credit cards*</u>
	Open credit card trades
<u>30 days late debt balances*</u>	Credit amount on open credit card trades
Total 30 days late debt balances	Open bankcard trades
Balance on 30 days late mortgage trades	Credit amount on revolving trades
<u>60 days late debt balances*</u>	<u>Number of HELOC loans*</u>
Total 60 days late debt balances	Open home equity line of credit trades
Balance on 60 days late mortgage trades	Home equity line of credit utilization ratio
<u>90+ days late debt balances*</u>	<u>Number of mortgages*</u>
Total 90 or more days delinquent debt balances	Fraction of mortgage to total debt
Total 90-180 days late debt balances	Mortgage trades
Balance on 90-180 days late mortgage trades	Open mortgage trades
Total derogatory debt balances	Mortgage type
Balance on derogatory mortgage trades	Open first mortgage trades
<u>Balance on installment loans*</u>	<u>Foreclosed first mortgages*</u>
Credit amount on open installment trades	Presently foreclosed first mortgages
Balance on open installment trades	Presently foreclosed first mortgage trades, last 24 months
<u>Balance on HELOC loans*</u>	<u>Auto loan*</u>
Credit amount on open home equity line of credit trades	Open auto loan trades
Balance on home equity line of credit trades	Monthly payment on open auto loan trades
Balance on open revolving trades	Balance on open auto loan trades
<u>Student debt*</u>	<u>Credit card debt*</u>
Balance on student trades	Balance on credit card trades
Credit amount on non-deferred student trades	Balance on revolving trades
	Balance on bankcard trades
<u>Worst status on any trades*</u>	<u>Trades ever 90 or more days delinquent or derogatory*</u>
Worst ever status on any trades in the last 24 months	90 days delinquencies in the last 6 months
Worst present status on any trades	90 days delinquencies in the last 12 months
	90 days delinquencies in the last 24 months
	90 days delinquencies in the last 36 months
<u>Worst status on credit cards*</u>	<u>Bankruptcy history*</u>
Worst present status on a revolving trade	Public record bankruptcies
Worst present status on bankcard trades	Chapter 7 bankruptcies
	Months since the most recent public record bankruptcy filed
<u>Fraction of 90 days late debt to total debt*</u>	
Worst present status on an open trade	
Fraction of 90 days delinquent debt to total debt	

## D Model Comparison

We trained a GBT with up to 10,000 trees, a learning rate of 0.3, and an early-stopping parameter of 1,000 to compare the performance of gradient boosting with deep neural networks. We also built on Table 10 by keeping our models' architecture the same, but expanded the training data by including observations up till the date specified by the training window. This exercise illustrates that while the performance of GBT remains similar, DNN benefits significantly from having more data to train on.

We also looked at the performance of the two models when we allow only the most recent 4 quarters for training.

We also investigated the SHAP values across four different models on the pooled sample: (1) logistic, (2) DNN, (3) GBT, and (4) the hybrid model. Table 22 summarizes these results.

Table 20: Model Comparison: DNN vs. GBT

Training Window*	Testing Window	AUC-score		Loss	
		DNN	GBT	DNN	GBT
2004Q1-2013Q4	2004Q1-2013Q4	0.9561	0.9429	0.2506	0.2832
2004Q1	2006Q1	0.9228	0.9241	0.3272	0.3236
2004Q2	2006Q2	0.9243	0.9251	0.3212	0.3190
2004Q3	2006Q3	0.9246	0.9264	0.3208	0.3162
2004Q4	2006Q4	0.9247	0.9256	0.3223	0.3194
2005Q1	2007Q1	0.9254	0.9261	0.3230	0.3209
2005Q2	2007Q2	0.9249	0.9258	0.3255	0.3228
2005Q3	2007Q3	0.9245	0.9253	0.3272	0.3250
2005Q4	2007Q4	0.9232	0.9243	0.3327	0.3284
2006Q1	2008Q1	0.9247	0.9249	0.3298	0.3289
2006Q2	2008Q2	0.9235	0.9244	0.3331	0.3306
2006Q3	2008Q3	0.9250	0.9253	0.3310	0.3284
2006Q4	2008Q4	0.9251	0.9256	0.3289	0.3281
2007Q1	2009Q1	0.9267	0.9275	0.3263	0.3233
2007Q2	2009Q2	0.9268	0.9276	0.3244	0.3222
2007Q3	2009Q3	0.9279	0.9288	0.3197	0.3183
2007Q4	2009Q4	0.9300	0.9307	0.3149	0.3133
2008Q1	2010Q1	0.9302	0.9308	0.3155	0.3141
2008Q2	2010Q2	0.9302	0.9308	0.3146	0.3128
2008Q3	2010Q3	0.9302	0.9308	0.3137	0.3117
2008Q4	2010Q4	0.9301	0.9307	0.3141	0.3124
2009Q1	2011Q1	0.9312	0.9317	0.3112	0.3097
2009Q2	2011Q2	0.9297	0.9302	0.3137	0.3126
2009Q3	2011Q3	0.9298	0.9302	0.3130	0.3117
2009Q4	2011Q4	0.9300	0.9302	0.3122	0.3115
2010Q1	2012Q1	0.9305	0.9307	0.3106	0.3099
2010Q2	2012Q2	0.9291	0.9293	0.3130	0.3124
2010Q3	2012Q3	0.9288	0.9290	0.3119	0.3113
2010Q4	2012Q4	0.9282	0.9286	0.3128	0.3121
2011Q1	2013Q1	0.9290	0.9293	0.3109	0.3101
2011Q2	2013Q2	0.9286	0.9289	0.3109	0.3103
2011Q3	2013Q3	0.9292	0.9294	0.3091	0.3085
2011Q4	2013Q4	0.9293	0.9295	0.3080	0.3074

Performance comparison of the two best performing machine learning classification models of consumer default risk. The model calibrations are specified by the training and testing windows. \* implies that all data was used up to the quarter specified. The results of predicted probabilities versus actual outcomes over the following 8Q (testing period) are used to calculate the loss metric and the AUC-score for 90+ days delinquencies within 8Q. DNN refers to deep neural network, GBT refers to gradient boosted trees. Source: Authors' calculations based on Experian Data.



Table 21: Model Comparison: DNN vs. GBT

Training Window Start	Training Window End	Testing Window	AUC-score		Loss	
			DNN	GBT	DNN	GBT
2004Q1	2004Q1	2006Q1	0.9229	0.9242	0.3273	0.3234
2004Q1	2004Q2	2006Q2	0.9239	0.9251	0.3226	0.3190
2004Q1	2004Q3	2006Q3	0.9251	0.9265	0.3209	0.3161
2004Q1	2004Q4	2006Q4	0.9244	0.9257	0.3218	0.3192
2004Q2	2005Q1	2007Q1	0.9253	0.9261	0.3236	0.3205
2004Q3	2005Q2	2007Q2	0.9245	0.9258	0.3252	0.3220
2004Q4	2005Q3	2007Q3	0.9244	0.9253	0.3271	0.3244
2005Q1	2005Q4	2007Q4	0.9237	0.9243	0.3282	0.3273
2005Q2	2006Q1	2008Q1	0.9238	0.9251	0.3300	0.3270
2005Q3	2006Q2	2008Q2	0.9237	0.9244	0.3313	0.3287
2005Q4	2006Q3	2008Q3	0.9247	0.9254	0.3286	0.3263
2006Q1	2006Q4	2008Q4	0.9249	0.9257	0.3277	0.3261
2006Q2	2007Q1	2009Q1	0.9266	0.9277	0.3246	0.3216
2006Q3	2007Q2	2009Q2	0.9268	0.9278	0.3235	0.3209
2006Q4	2007Q3	2009Q3	0.9278	0.9288	0.3196	0.3172
2007Q1	2007Q4	2009Q4	0.9297	0.9307	0.3149	0.3124
2007Q2	2008Q1	2010Q1	0.9293	0.9306	0.3169	0.3133
2007Q3	2008Q2	2010Q2	0.9292	0.9304	0.3164	0.3129
2007Q4	2008Q3	2010Q3	0.9288	0.9302	0.3163	0.3128
2008Q1	2008Q4	2010Q4	0.9285	0.9299	0.3179	0.3142
2008Q2	2009Q1	2011Q1	0.9292	0.9307	0.3157	0.3121
2008Q3	2009Q2	2011Q2	0.9276	0.9291	0.3189	0.3150
2008Q4	2009Q3	2011Q3	0.9276	0.9291	0.3185	0.3147
2009Q1	2009Q4	2011Q4	0.9276	0.9294	0.3180	0.3138
2009Q2	2010Q1	2012Q1	0.9291	0.9302	0.3144	0.3113
2009Q3	2010Q2	2012Q2	0.9277	0.9290	0.3164	0.3133
2009Q4	2010Q3	2012Q3	0.9279	0.9288	0.3140	0.3119
2010Q1	2010Q4	2012Q4	0.9273	0.9284	0.3154	0.3132
2010Q2	2011Q1	2013Q1	0.9281	0.9291	0.3137	0.3111
2010Q3	2011Q2	2013Q2	0.9275	0.9286	0.3142	0.3112
2010Q4	2011Q3	2013Q3	0.9280	0.9292	0.3127	0.3095
2011Q1	2011Q4	2013Q4	0.9282	0.9293	0.3116	0.3084
Average			0.9267	0.9278	0.3202	0.3172

Performance comparison of the two best performing machine learning classification models of consumer default risk. The model calibrations are specified by the training and testing windows. The results of predicted probabilities versus actual outcomes over the following 8Q (testing period) are used to calculate the loss metric and the AUC-score for 90+ days delinquencies within 8Q. DNN refers to deep neural network, GBT refers to gradient boosted trees. Source: Authors' calculations based on Experian Data.

Table 22: Shap Values across Models

Feature	Hybrid	Logistic	DNN	GBT
Worst status on any trades*	0.596 (1)	0.095 (1)	0.155 (1)	1.038 (1)
Months since the oldest trade was opened	0.181 (2)	0.036 (4)	0.077 (4)	0.303 (4)
Months since the most recent 90 or more days delinquency	0.177 (3)	0.018 (9)	0.068 (5)	0.329 (2)
Number of collections*	0.164 (4)	0.056 (2)	0.132 (3)	0.211 (9)
Number of open credit cards*	0.163 (5)	0.039 (3)	0.137 (2)	0.245 (6)
Total debt balances*	0.161 (6)	0.01 (13)	0.028 (10)	0.308 (3)
Credit card utilization ratio	0.144 (7)	0.002 (43)	0.015 (23)	0.275 (5)
90+ days late debt balances*	0.125 (8)	0.016 (11)	0.033 (7)	0.221 (8)
Credit amount on unsatisfied derogatory trades	0.117 (9)	0.001 (57)	0.018 (18)	0.221 (7)
Months since the most recent 30-180 days delinquency	0.101 (10)	0.003 (36)	0.045 (6)	0.175 (11)
Inquiries made in the last 12 months	0.097 (11)	0.018 (8)	0.023 (12)	0.174 (12)
Monthly payment on all debt	0.091 (12)	0.0 (88)	0.008 (45)	0.179 (10)
Utilization ratio	0.089 (13)	0.008 (18)	0.017 (20)	0.165 (13)
Months since the most recently opened credit card trade	0.07 (14)	0.01 (14)	0.023 (13)	0.127 (14)
Credit card debt*	0.066 (15)	0.018 (10)	0.02 (16)	0.125 (15)
Balance on collections	0.063 (17)	0.003 (35)	0.008 (41)	0.123 (17)
Balance on installment loans*	0.063 (16)	0.001 (54)	0.014 (25)	0.124 (16)
Credit amount paid down on open first mortgage trades	0.062 (18)	0.001 (51)	0.013 (31)	0.114 (19)
Months since the most recently opened first mortgage	0.057 (19)	0.001 (65)	0.015 (22)	0.114 (18)
Monthly payment on credit card trades	0.052 (20)	0.009 (16)	0.014 (30)	0.097 (22)
Auto loan*	0.051 (21)	0.003 (29)	0.017 (19)	0.1 (21)
Monthly payment on open first mortgage trades	0.05 (22)	0.001 (74)	0.006 (54)	0.1 (20)
Number of mortgages*	0.045 (23)	0.028 (6)	0.029 (9)	0.079 (26)
Balance on HELOC loans*	0.044 (25)	0.006 (22)	0.006 (53)	0.087 (23)
Installment utilization ratio	0.044 (24)	0.004 (27)	0.013 (33)	0.082 (24)
Student debt*	0.042 (26)	0.008 (17)	0.014 (26)	0.079 (25)
Months since the most recently opened auto loan trade	0.039 (27)	0.001 (67)	0.022 (14)	0.072 (28)
Charge-off amount on unsatisfied charge-off trades	0.038 (28)	0.001 (66)	0.002 (69)	0.076 (27)
Fraction of 90 days late debt to total debt*	0.035 (29)	0.029 (5)	0.024 (11)	0.051 (36)
Worst ever status on a credit card trade in the last 24 months	0.034 (30)	0.001 (55)	0.014 (24)	0.061 (30)

This table reports the Shapley values for four selected machine learning classification models of consumer default risk. We sorted the features based on the feature’s relative rank (in parentheses) using the hybrid model. Source: Authors’ calculations based on Experian Data.

Table 23: Neural networks comparison: Loss & Accuracy

Model	In-sample Loss		Out-of-sample Loss	
	w/o Dropout	Dropout	w/o Dropout	Dropout
Logistic Regression	0.3451	0.3451	0.3449	0.3450
1 layer	0.3109	0.3106	0.3122	0.3116
2 layers	0.2965	0.2900	0.3078	0.3003
3 layers	0.2804	0.2460	0.3047	0.2744
4 layers	0.2669	0.2142	0.3005	0.2575
5 layers	0.2534	0.2013	0.2978	0.2506

Model	In-sample Accuracy		Out-of-sample Accuracy	
	w/o Dropout	Dropout	w/o Dropout	Dropout
Logistic Regression	0.8564	0.8564	0.8566	0.8566
1 layer	0.8687	0.8688	0.8681	0.8684
2 layers	0.8751	0.8787	0.8705	0.8736
3 layers	0.8829	0.9017	0.8729	0.8862
4 layers	0.8897	0.9163	0.8755	0.8943
5 layers	0.8968	0.9230	0.8785	0.8981

In-sample and out-of-sample loss (categorical cross-entropy) and accuracy for neural networks of different depth and for logistic regression. Models are calibrated and evaluated on the pooled sample (2004Q1 - 2013Q4). Source: Authors' calculations based on Experian Data.

## E Comparison with Credit Scores

The credit score is a summary indicator intended to predict the risk of default by the borrower and it is widely used by the financial industry. For most unsecured debt, lenders typically verify a perspective borrower's credit score at the time of application and sometimes a short recent sample of their credit history. For larger unsecured debts, lenders also typically require some form of income verification, as they do for secured debts, such as mortgages and auto loans. Still, the credit score is often a key determinant of crucial terms of the borrowing contract, such as the interest rate, the downpayment or the credit limit.

The most widely known credit score is the FICO score, a measure generated by the Fair Isaac Corporation, which has been in existence in its current form since 1989. Each of the three major credit reporting bureaus— Equifax, Experian and TransUnion— also have their own proprietary credit scores. Credit scoring models are not public, though they are restricted by the law, mainly the Fair Credit Reporting Act of 1970 and the Consumer Credit Reporting Reform Act of 1996. The legislation mandates that consumers be made aware of the 4 main factors that may affect their credit score. Based on available descriptive materials from FICO and the credit bureaus, these are payment history and outstanding debt, which account for more than 60% of the variation in credit scores, followed by credit history, or the age of existing accounts, which explains 15-20% of the variation, followed by new accounts and types of credit used (10-5%) and new "hard" inquiries, that is credit report inquiries coming from perspective lenders after a borrower initiated credit application.

U.S. law prohibits credit scoring models from considering a borrower's race, color, religion, national origin, sex and marital status, age, address, as well as any receipt of public assistance, or the exercise of any consumer right under the Consumer Credit Protection Act. The credit score cannot be based on information not found in a borrower's credit report, such as salary, occupation, title, employer, date employed or employment history, or interest rates being charged on particular accounts. Finally, any items in the credit report reported as child/family support obligations are not permitted, as well as "soft" inquiries<sup>29</sup> and any information that is not proven to be predictive of future credit performance.

---

<sup>29</sup>These include "consumer-initiated" inquiries, such as requests to view one's own credit report, "promotional inquiries," requests made by lenders in order to make pre-approved credit offers, or "administrative inquiries," requests made by lenders to review open accounts. Requests that are marked as coming from employers are also not counted.




Figure 13: Histogram of the credit score in our data by year for selected years. Source: Source: Authors' calculations based on Experian data.





Figure 14: Absolute value of rank correlation with realized default rate for the credit score and model predicted default probability (a) and Gini coefficients for the credit score and model predicted default probability by quarter (b). Source: Authors' calculations based on Experian data.



(a) Rank Correlation with Realized Default Rates

(b) Gini Coefficients by Quarter

Figure 15: Absolute value of rank correlation with realized default rate for the credit score and model predicted default probability (a) and Gini coefficients for the credit score and model predicted default probability by quarter (b) for the population of current borrowers. Borrowers are current if they have not outstanding delinquencies. Source: Authors' calculations based on Experian data.

## F Cost Savings for Consumers

Table 24: Distribution of Customers by Credit Score and Predicted Default

		Credit Score			
		Subprime, Near Prime	Prime Low	Prime Mid	Prime High, Superprime
Predicted Default Bin	1	36.22%	3.63%	1.24%	0.41%
	2	3.87%	3.41%	2.41%	1.08%
	3	1.09%	2.52%	4.01%	3.49%
	4	0.33%	1.21%	3.44%	31.64%

This table reports the share of customers in each predicted credit score categories and corresponding predicted default probability bins. Customers are classified based on credit scores and predicted default probabilities at account origination for each of their credit cards included in the balances. Source: Authors' calculations based on Experian data.

Table 25: Cost Savings for Consumers: Descriptive Statistics

		Credit Score				Credit Score	
		Prime Low	Prime Mid	Prime High	Superprime		
Subprime, Near Prime	Prime	Prime Low	Prime Mid	Prime High	Superprime		
<b>Age</b>							
<b>PP bin</b>	1	42	44	47	47	51	161
	2	40	41	44	44	50	160
	3	43	42	42	42	47	193
	4	47	47	45	45	53	230
<b>Household Income (\$)</b>							
<b>PP bin</b>	1	57,673	65,929	74,326	74,326	91,429	0.14
	2	65,291	66,527	74,474	74,474	92,751	0.27
	3	83,843	77,536	75,962	75,962	93,328	0.37
	4	107,025	101,535	97,465	97,465	109,494	0.46
<b>Total Debt Balances (\$)</b>							
<b>PP bin</b>	1	65,275	84,746	91,062	91,062	105,104	7,401
	2	105,681	94,048	103,516	103,516	124,871	7,015
	3	174,560	126,318	106,815	106,815	125,445	8,763
	4	260,262	200,065	163,614	163,614	103,639	10,021
<b>Total bankcard balances (\$)</b>							
<b>PP bin</b>	1	3,822	4,332	3,968	3,968	3,987	4,442
	2	10,480	6,955	5,502	5,502	4,556	11,383
	3	16,330	10,606	6,794	6,794	4,968	17,104
	4	20,660	14,464	9,824	9,824	4,200	21,074
<b>Credit Card Limit (\$)</b>							
<b>PP bin</b>	1	10,666	13,621	18,930	18,930	29,238	1.17
	2	20,740	18,832	22,303	22,303	31,782	0.93
	3	34,333	28,705	26,571	26,571	35,260	0.80
	4	49,673	44,007	40,584	40,584	42,647	0.64
<b>90+ DPD on All Accounts, Next 24 Months</b>							
<b>PP bin</b>	1	0.62	0.44	0.37	0.37	0.35	0.22
	2	0.36	0.29	0.25	0.25	0.21	0.15
	3	0.29	0.23	0.18	0.18	0.14	0.13
	4	0.24	0.19	0.13	0.13	0.08	0.11
<b>90+ DPD on Bankcards, Next 24 Months</b>							
<b>PP bin</b>	1	0.35	0.37	0.37	0.37	0.35	0.22
	2	0.21	0.25	0.25	0.25	0.21	0.15
	3	0.14	0.18	0.18	0.18	0.14	0.13
	4	0.08	0.13	0.13	0.13	0.08	0.11
<b>Credit Score, Current</b>							
<b>PP bin</b>	1	606	661	694	694	733	0.60
	2	644	671	701	701	739	0.34
	3	659	684	711	711	746	0.26
	4	673	697	724	724	775	0.20
<b>Predicted Probability, Current</b>							
<b>PP bin</b>	1	0.37	0.44	0.37	0.37	0.35	0.22
	2	0.25	0.25	0.25	0.25	0.21	0.15
	3	0.18	0.18	0.18	0.18	0.14	0.13
	4	0.13	0.13	0.13	0.13	0.08	0.11

This table reports the descriptive statistics on customers based on their credit score and predicted default probability based on our hybrid model. Credit score categories and predicted default probability bins are computed at account origination for credit cards. Source: Authors' calculations based on Experian data.