

# Normalizing Pandemic Data for Credit Scoring

Joseph L Breeden  
Deep Future Analytics LLC  
1600 Lena St., Suite E3,  
Santa Fe, NM 87505 USA  
breeden@deepfutureanalytics.com

March 23, 2025

## Abstract

The COVID-19 pandemic created abnormal credit risk conditions that did not align well with pre-2020 credit scores. Since the pandemic, most organizations have either excluded the period 2020-2021 from their modeling or included it without adjustment, leaving it as noise in the data. Model validators and examiners have been divided about requiring one of these approaches or defaulting to model developer judgment. None of this is ideal from a model development perspective.

We have found that a technical solution is available. Our analysis uses lifecycle and environment outputs from an Age-Period-Cohort analysis as fixed offsets to the credit score development. Panel data is used, so the credit score is developed with a discrete time survival model approach. We tested logistic regression and stochastic gradient boosted regression trees as estimators with the panel data and APC inputs.

For this research, we used Fannie Mae data. The APC model was estimated on the full available history, from 2005 through 2024. The origination scores were estimated on two-year periods from 2016 through 2024 and tested on all other periods, including a score that was developed on the full period. All models were also tested on comparably prepared data from Freddie Mac for cross-validation.

# 1 Introduction

The COVID-19 pandemic created unprecedented economic conditions that fundamentally disrupted traditional credit risk assessment. According to research on marketplace lending, the probability of loan default increased from 0.056 in the pre-pandemic period to 0.079 during the pandemic. This significant shift occurred despite massive government intervention programs designed to stabilize the economy [6].

These intervention programs, while beneficial for short-term economic stability, created additional challenges for credit risk assessment. A study by [3] found that default intensities shifted from long-range to short-range dependence during the COVID-19 period, making historical credit performance much less relevant for credit prediction. This phenomenon contrasts sharply with previous financial crises, where long-memory patterns remained relatively stable. The current research appears to present a solution to this challenge.

Financial institutions faced significant challenges in implementing credit risk models during this period. The disconnect between economic conditions and credit performance led many institutions to replace advanced credit scores with simple bureau score-based decision trees.

Traditional credit scoring models faced particular challenges during the pandemic for several reasons. The rapid implementation of forbearance programs artificially suppressed delinquencies, creating a disconnect between observed credit performance and actual credit risk. Consumer spending patterns changed dramatically as lockdowns limited discretionary spending opportunities, leading to decreased credit utilization and paradoxically improved credit scores during a period of economic distress. This created what Breeden [8] termed “Macroeconomic Adverse Selection”, where the macroeconomic environment significantly influenced the distribution of consumers applying for loans in ways not reflected in their credit scores. The impact of the pandemic was heterogeneous across different segments of borrowers, with certain industries and geographic regions experiencing more severe effects [30].

A fundamental limitation of traditional credit scoring approaches is their reliance on cross-sectional data rather than panel data. Cross-sectional models provide only a snapshot at a single point in time, failing to capture the dynamic nature of credit risk as it changes through a loan’s lifecycle and across different economic environments.

Panel data approaches, which track the same borrowers over multiple time periods, offer several advantages for credit risk modeling [36]. By knowing when a loan defaults, the model can be designed to distinguish between defaults that occur due to the maturing of the loan, shifts in the economic environment, or the intrinsic credit risk of the loan. During economic shocks like the pandemic, these advantages become even more critical as they allow models to adapt to changing conditions and capture time-varying effects that cross-sectional approaches cannot address.

This paper introduces a method to normalize pandemic data for use in credit scoring. The lifecycle and environment functions from an Age-Period-Cohort (APC) model are used to normalize panel-structured performance data. With this normalization, logistic regression and stochastic gradient boosted trees are created on Fannie Mae and Freddie Mac mortgage data to create origination scores that maintain predictive power despite the pandemic’s unprecedented conditions.

## 2 Model Architecture

Many lending analysts have tried to create macroeconomic indices to measure the actual impact of the pandemic on consumers [16, 21, 20, 29]. The author tried creating an index of “persons receiving income”, which would include employed people and those receiving government support. Although this and other indices made sense, none could capture the impacts of the various forbearance programs and benefits. Rather than a macroeconomic-based approach to normalization, the most effective method was to measure the net impact via age-period cohort modeling and then adjust the score estimation to use the observed environmental impacts for normalization.

### 2.1 Fundamentals of APC Modeling

Age-Period-Cohort (APC) analysis provides an effective framework for modeling default risk over time [38, 26]. Originally developed in epidemiology and sociology [28], APC models have been used in loss forecasting and stress testing [7]. These models decompose credit performance into three fundamental components: lifecycle effects versus age of the account, environmental impacts versus calendar date, and credit quality by vintage.

The basic APC model can be expressed as

$$\text{log-odds}PD(a, t, v) = F(a) + G(v) + H(t) \quad (1)$$

$PD$  is defined in this context as Ever 60+ DPD Rate, meaning a "default" is recorded on the first occurrence of an account becoming 60 or more days past due.  $a$  is the age of the loan,  $t$  is the calendar time,  $v$  is the vintage date,  $F(a)$  represents the lifecycle function,  $H(t)$  represents the environment function, and  $G(v)$  represents the vintage quality function.

This decomposition allows researchers to separate the effects of vintage quality from loan aging and macroeconomic conditions. During economic shocks like the COVID-19 pandemic, the environment function captures the impact of changing consumer financial conditions, while the vintage function reflects changes in borrower quality. The ability to measure actual impacts via  $H(t)$ , rather than trying to explain calendar date impacts purely from macroeconomic factors, was a key to success during the pandemic.

One challenge in APC modeling is the identification problem arising from the linear relationship  $a = t - v$  [27]. Various approaches have been proposed to address this issue, including imposing constraints on the functions. In credit risk applications, a common approach is to represent the model as

$$\log(PD(a, t, v)) = b_0 + b_1a + b_2v + F'(a) + G'(v) + H'(t) \quad (2)$$

where  $F'(a)$ ,  $G'(v)$ , and  $H'(t)$  are nonlinear functions with zero mean and no linear component. This formulation resolves the linear trend ambiguity. The analyst must then decide how best to allocate the trend, considering the details of the problem. For the analysis here, the environment function is chosen to have zero trend, which is most appropriate with long time histories spanning more than one recession. If less than one economic cycle had been available, fitting to economic factors while allowing for a secular trend can resolve the environmental trend uncertainty.

The functions are estimated using a Bayesian Age-Period-Cohort algorithm [32], in order to constrain the estimates for recent vintages and older dates where few observations are available.

## 2.2 Panel Logistic Regression with APC Inputs

Many credit scoring techniques exist which leverage survival modeling principles [22, 34, 33, 5]. The most common of these is Cox proportional hazards

[18, 35, 31], which is widely used, but suffers from instability in a credit scoring applications [11].

To implement discrete time survival models with APC inputs, we use panel logistic regression with an augmented model matrix [14]. One row is created per account per month active, and no more than one entry for a default event. Censoring the data after default means implicitly that the denominator for  $PD$  will be the previous month's open accounts. This structure allows the analyst to include a parallel model of the probability of prepayment in order to consider competing risks.

The model matrix becomes

$$X_i = \begin{bmatrix} 0 & F(1) & H(t_1) & s_{i1} & s_{i2} & \cdots & s_{ik} \\ 0 & F(2) & H(t_2) & s_{i1} & s_{i2} & \cdots & s_{ik} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & F(a_i) & H(t_i) & s_{i1} & s_{i2} & \cdots & s_{ik} \end{bmatrix} \quad (3)$$

where the first column is the default indicator, the second column is the lifecycle at each age, the third column is the environment function at each time, and the remaining columns are scoring covariates. During model estimation, the coefficients for the lifecycle and environment columns are fixed at 1.0.

By including the APC lifecycle and environment functions as fixed inputs, the panel logistic regression model can be expressed as:

$$\log \left( \frac{PD(a, v, t, i)}{1 - PD(a, v, t, i)} \right) = F(a) + H(t) + \sum_j c_j s_{ij} + \sum_v \beta_v \delta(v) + \varepsilon(a, v, t, i) \quad (4)$$

where  $s_{ij}$  are the available attributes at origination for account  $i$ , and the  $c_j$  are the coefficients to be estimated. This approach allows  $F(a) + H(t)$  to specify the mean of the distribution at each forecast month, with the logistic regression model estimating the distribution of account risk centered about that mean. In other words, the credit score models the residual risk after adjusting for lifecycle and environment, thus normalizing the performance data for account age and the pandemic conditions.

This approach significantly improves out-of-time performance compared to traditional credit scoring methods [14]. By providing lifecycle and environment functions as fixed inputs to the model, we can achieve more stable coefficients for risk factors and better capture the impact of changing economic conditions.

### 2.3 Stochastic Gradient Boosted Trees

Machine learning methods, particularly gradient boosted trees, have gained popularity in credit scoring due to their ability to capture complex, non-linear relationships in data. Stochastic Gradient Boosted Regression Trees (SGBT) combines bagging with gradient boosting to create an ensemble of ensembles of trees [19, 25].

The basic idea of gradient boosting is to build subsequent models on the residuals of previous models, computing the gradient of a fitness function to provide weights to each model trained. Stochastic gradient boosting adds randomization to this process, building different gradient boosted ensembles for each data subsample.

$$F_K(x) = c + \sum_{k=1}^K \nu h_k(x) \quad (5)$$

where each  $h_k(x)$  is a model and  $\nu$  is a learning rate parameter. This approach can significantly reduce computation times while maintaining acceptable accuracy.

When regression models are used for the  $h_k(x)$ , then lifecycle and environment may be included as fixed offsets, just as was done for the discrete time survival model [14]. The input data is again structured as a panel, so the resulting forecasts are monthly, conditional  $PD$ . Other researchers have developed combinations of survival modeling and stochastic gradient boosted trees [17, 4, 37, 1, 2], although for the current problem, we prefer the APC approach of incorporating an empirical environment function for normalization of the credit score development.

## 3 Fannie Mae and Freddie Mac Data for Origination Scores

To demonstrate the approach, loan-level data from Fannie Mae and Freddie Mac were analyzed. The Freddie Mac Single Family Loan-Level Dataset covers approximately 53.8 million mortgages originated between January 1, 1999, and June 30, 2024, with monthly loan performance data through September 30, 2024 [24]. Comparable data from the Fannie Mae Single-Family Loan Performance Data was prepared for modeling [23]:

For our analysis, we focus on loans with FICO scores between 660 and 780, following the approach described in the search results. This creates a dataset with 622,452 unique loans of which 4,346 defaulted for a lifetime default rate of 0.7%. The data includes vintages from January 1999 through November 2019, with performance data from January 2017 through December 2019.

The origination scoring variables available from both Fannie Mae and Freddie Mac are given in Table 1.

Table 1: Variables for credit scoring

Variable Name
Borrower Credit Score At Origination
Loan Purpose
Number Of Borrowers
Original Combined Loan To Value Ratio (CLTV)
Original Debt To Income Ratio (DTI)
Original Interest Rate
Property Type

This set of variables is sufficient to demonstrate the approach, but not to highlight the nonlinear modeling capabilities of the stochastic gradient boosted regression trees.

To prepare the data for our panel modeling approach, we create an augmented data matrix where each loan has multiple observations corresponding to different points in its lifecycle. For each observation, we include both static variables (such as bureau credit score and LTV at origination) and the values of the APC functions.

## 4 Results

### 4.1 Age-Period-Cohort Model

For the APC analysis, the full Fannie Mae data was analyzed. This was necessary in order to estimate the environment function over the full date range. For all in-sample and out-of-sample test results, the actual environment function was used. This means that out-of-sample or out-of-time refers

to the data used for training the score. This approach is called an "ideal scenario validation", because we do not want to confuse our scoring model tests with errors from predicting the environment function using macroeconomic variables.

The environment function clearly shows the 2001 and 2009 recessions and the 2020 pandemic, Figure 2. The lifecycle functions, Figure 1, mirror those observed previously [9].

The APC vintage functions, Figure 3, clearly show the strong periods of adverse selection from 2006 through 2008 [12] and more recently from 2022 through 2024 [10]. These periods have previously been explained as arising from sudden shifts in interest rates, causing a shift in the intrinsic risk of the loan applicants in ways not visible from standard scoring attributes. This effect was also observed in auto loans [13] and other consumer loans [15].

For purposes of demonstrating normalizing pandemic data for scoring, the vintage function was not included in the credit score, and macroeconomic adverse selection was not estimated after score construction. If these models were put into production, researchers should continuously track vintage residuals and include these in the forecast in order to align the credit score forecasts from different periods. Vintage residuals are not exactly the same as the APC vintage function. Conceptually, they are equivalent to subtracting the credit score aggregated by vintage from the vintage function.

## 4.2 Credit Scores

To test the credit scoring approach, models were estimated on two-year periods and tested on all other two-year periods. For each type of credit score, this process created four separate models plus a fifth model spanning the entire range.

A panel logistic regression model was created with lifecycle and environment as fixed offsets, as discussed earlier. All of the scoring variables from Table 1 were significant and retained in all models. Table 2 shows the Gini coefficients for each model. In-sample test values are shown in bold. Out-of-time test values are shown in regular type. The table is designed to facilitate the comparison of panel logistic regression models with and without the APC offsets. As shown in the Difference section, the difference is insignificant except for the period 2020-2021, in which case the model with an offset had a Gini coefficient on average 0.28 higher than the panel regression without a coefficient. The model built on the entire span had a comparable gain,

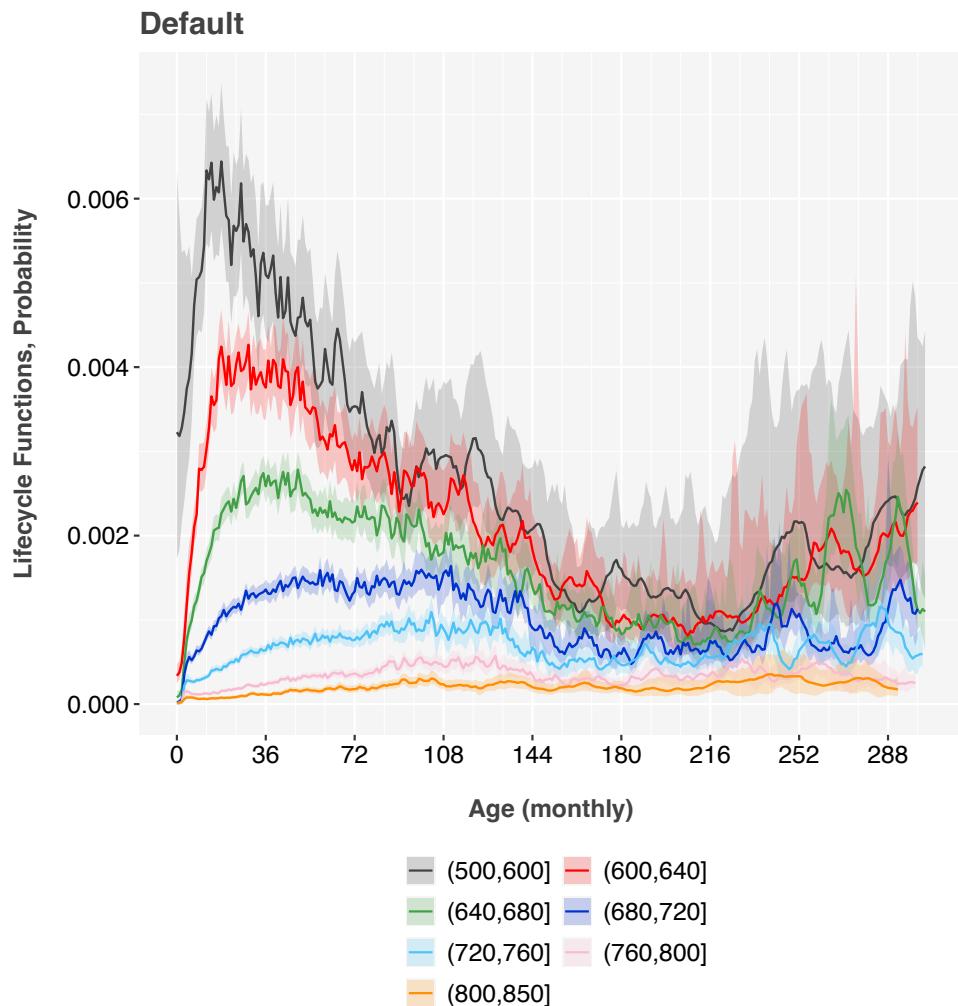


Figure 1: APC lifecycles obtained from Fannie Mae data, segmented by bureau score bands. The y-axis is in units of monthly probability of default.

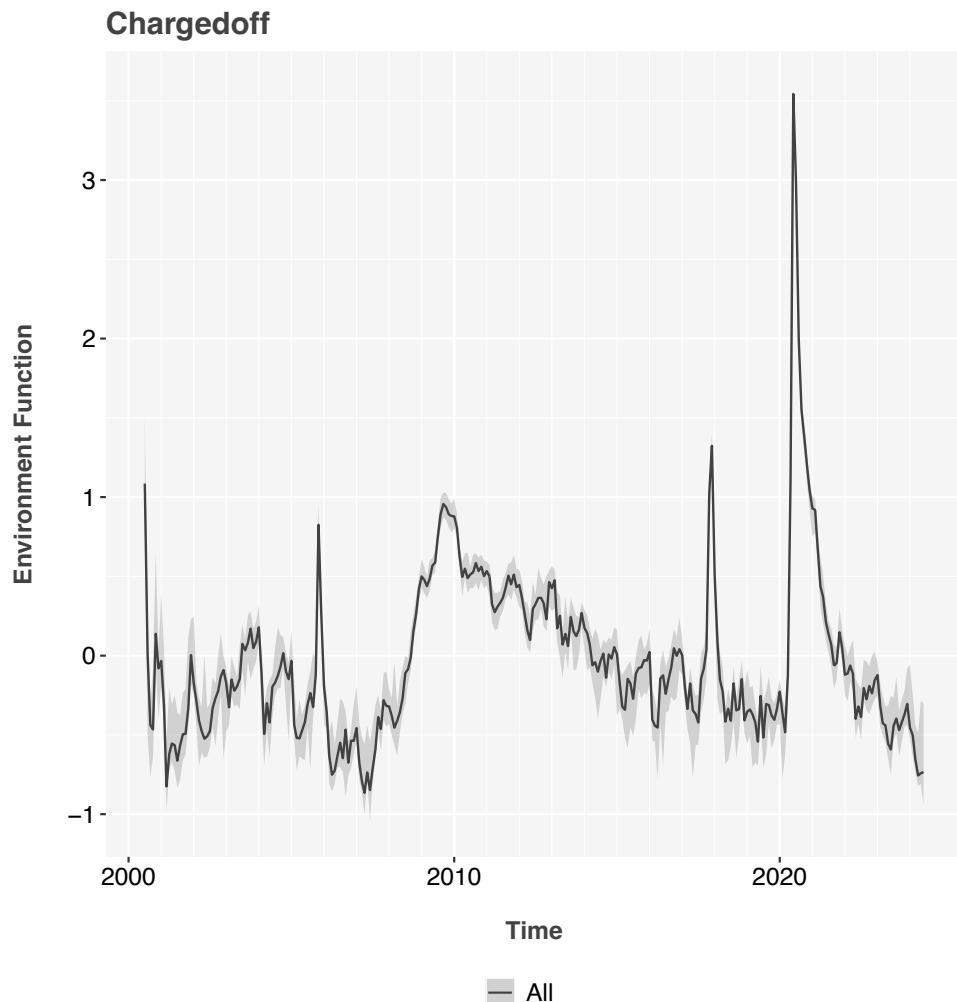


Figure 2: APC environment function for Fannie Mae data. Testing showed that the function was equivalent across segments considering confidence intervals. Therefore, a single overall function was used in later analysis. The y-axis is in units of log-odds of default.

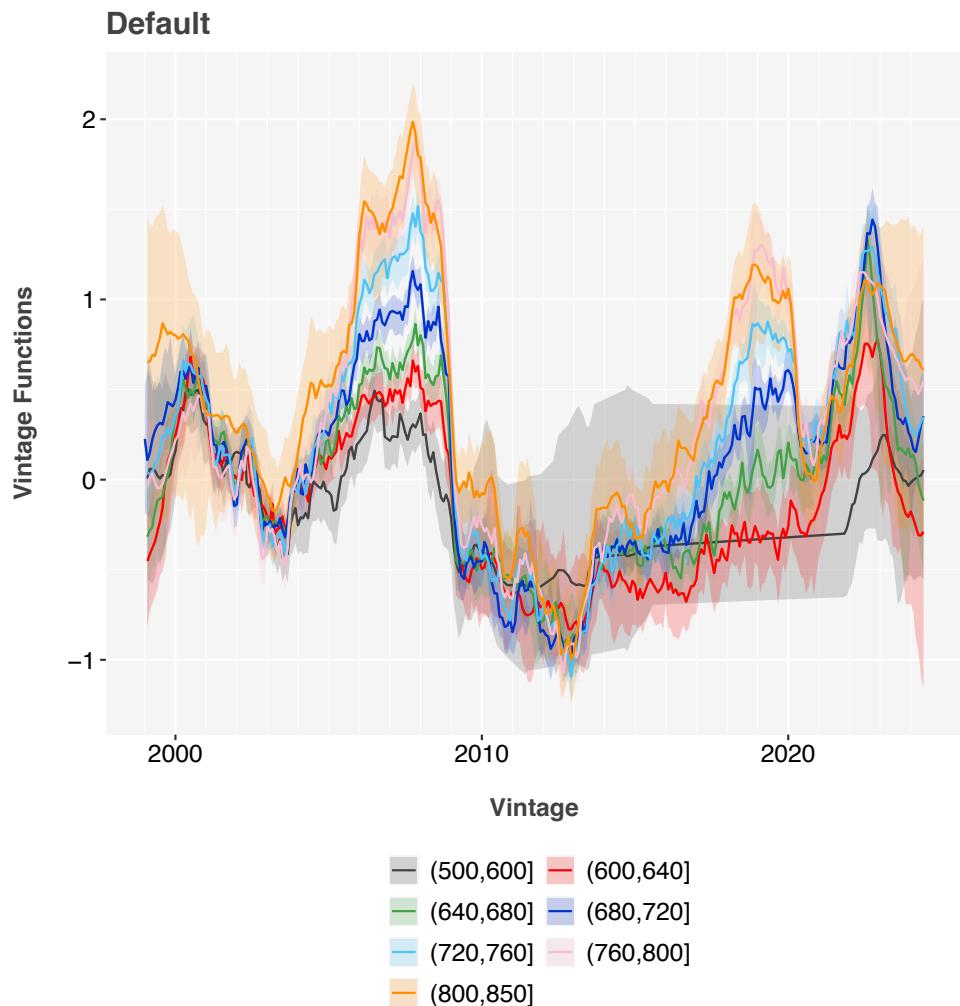


Figure 3: APC vintage functions, segmented by score band. The y-axis is in units of log-odds of default, so higher values correspond to riskier vintages even after segmenting by bureau score.

presumably because of the inclusion of the 2020-2021 period.

The section for Freddie Mac was included as a fully out-of-sample test, since none of this data was included in the model. The comparison of Freddie Mac to Fannie Mae could give an idea as to how much a score would degrade if built on Fannie Mae data and then applied to a lender's in-house loan origination. The Freddie Mac results generally show a reduction in Gini of 0.08 relative to Fannie Mae, but with the same pattern regarding the 2020-2021 period.

For the second test, a panel stochastic gradient boosted regression tree (SGBRT) was created. The SGBRT models were estimated on 60% of the in-sample data with 40% as a cross-validation hold-out sample. The meta-parameters for the SGBRT algorithm (number of leaves, tree depth, learning rate, number of trees) were optimized in-sample on the 2016-2017 time slice and used throughout. One model was given APC lifecycle and environment as fixed offsets, and the other was not. Those results are shown in Table 3 with the same format as for the panel logistic regression outputs.

The panel SGBRT models show slightly better performance than panel logistic regression in all tests, but the difference is not significant. The pattern across test periods is again the same as for Table 2, with the 2020-2021 period showing the greatest when including the APC inputs. The Freddie Mac out-of-sample testing was even closer to the Fannie Mae results.

Previous studies have shown improvement when using APC inputs, in-sample and most notably out-of-sample. The mortgage data is different, because it is a very large data set with a wide range of vintages. When modeled over a short, two-year time window in a calm economic period with a limited set of available scoring variables, the advantages of APC inputs largely disappear. For smaller, more dynamic portfolios of shorter-term loans, the APC lifecycle will be an important contributor.

During very dynamic economic periods, as seen in 2020-2021, the APC environment input is essential to creating an effective model. Perhaps ironically, the model built during this time period with APC inputs performs better in all time periods than the models built in-sample on those other time periods. This is likely due to the increased ability to separate environmental effects from scoring attribute dynamics when the environment is dynamic and clearly identified.

The lack of significant improvement from SGBRT as compared to LR is not a surprise. Much has been written trying to determine which modeling algorithm is best without paying sufficient attention to the data being

Table 2: Logistic regression comparisons. Bold values are in-sample, in-time. All Freddie Mac outputs are out-of-sample.

		Fannie Mae		Test Period												Difference		
				No Offset						With Offset								
				2016-2017	2018-2019	2020-2021	2022-2023	2016-2024	2016-2017	2018-2019	2020-2021	2022-2023	2016-2024	2016-2017	2018-2019	2020-2021	2022-2023	2016-2024
Train Period	2016-2017	<b>0.60</b>	0.59	0.46	0.54	0.49	<b>0.65</b>	0.59	0.73	0.52	0.70	0.55	0.00	0.28	-0.02	0.21		
	2018-2019	0.60	<b>0.59</b>	0.45	0.55	0.49	0.60	<b>0.60</b>	0.73	0.52	0.70	0.55	0.00	0.28	-0.03	0.22		
	2020-2021	0.58	0.56	<b>0.47</b>	0.51	0.49	0.53	<b>0.74</b>	0.50	0.69	0.50	0.01	0.27	0.00	0.21			
	2022-2023	0.56	0.56	0.42	<b>0.56</b>	0.48	0.57	0.54	0.74	0.52	<b>0.70</b>	0.03	0.31	-0.02	0.23			
	2016-2024	0.58	0.57	0.45	0.55	<b>0.49</b>	0.62	0.56	0.74	0.52	<b>0.70</b>	0.03	-0.01	0.28	-0.03	0.21		
Freddie Mac				No Offset						With Offset						Difference		
				2016-2017	2018-2019	2020-2021	2022-2023	2016-2024	2016-2017	2018-2019	2020-2021	2022-2023	2016-2024	2016-2017	2018-2019	2020-2021	2022-2023	2016-2024
2016-2017	0.48	0.49	0.41	0.48	0.42	0.53	0.49	0.64	0.46	0.59	0.45	0.00	0.23	-0.019	0.18			
2018-2019	0.48	0.49	0.40	0.50	0.42	0.53	0.50	0.64	0.47	0.59	0.45	0.012	0.24	-0.023	0.17			
2020-2021	0.46	0.45	0.42	0.45	0.41	0.46	0.42	0.64	0.45	0.58	0.55	-0.030	0.22	0.001	0.17			
Train Period	2022-2023	0.44	0.46	0.37	0.51	0.40	0.50	0.47	0.65	0.49	0.60	0.59	0.010	0.27	-0.019	0.20		
	2016-2024	0.46	0.47	0.40	0.49	0.42	0.50	0.46	0.65	0.47	0.59	0.60	-0.005	0.24	-0.023	0.18		
Source Comparison				No Offset						With Offset						Difference		
				2016-2017	2018-2019	2020-2021	2022-2023	2016-2024	2016-2017	2018-2019	2020-2021	2022-2023	2016-2024	2016-2017	2018-2019	2020-2021	2022-2023	2016-2024
2016-2017	0.12	0.10	0.045	0.057	0.074	0.12	0.098	0.098	0.088	0.051	0.11							
2018-2019	0.12	0.10	0.044	0.054	0.072	0.12	0.097	0.087	0.067	0.050	0.11							
2020-2021	0.12	0.11	0.051	0.063	0.079	0.13	0.11	0.096	0.057	0.12	0.11							
Train Period	2022-2023	0.12	0.10	0.048	0.051	0.074	0.12	0.098	0.088	0.068	0.052	0.11						
	2016-2024	0.12	0.10	0.051	0.057	0.077	0.11	0.098	0.091	0.054	0.054	0.11						

Table 3: Stochastic gradient boosted regression tree comparisons. Bold values are in-sample. All Freddie Mac outputs are out-of-sample.

Fannie Mae		Test Period										Difference				
		No Offset					With Offset									
		2016-2017	2018-2019	2020-2021	2022-2023	2016-2024	2016-2017	2018-2019	2020-2021	2022-2023	2016-2024	2016-2017	2018-2019	2020-2021	2022-2023	2016-2024
Train Period	2016-2017	<b>0.62</b>	0.59	0.45	0.54	0.48	<b>0.67</b>	0.61	0.73	0.52	0.70	0.05	0.01	0.28	-0.02	0.22
2018-2019	0.60	<b>0.61</b>	0.45	0.54	0.49	0.66	<b>0.62</b>	0.73	0.55	0.71	0.06	0.01	0.28	0.01	0.23	0.23
2020-2021	0.53	0.53	<b>0.52</b>	0.42	0.49	0.61	<b>0.55</b>	<b>0.75</b>	0.52	0.71	0.08	0.02	0.23	0.10	0.22	0.23
2022-2023	0.57	0.58	0.43	<b>0.58</b>	0.49	0.64	<b>0.52</b>	0.74	0.59	<b>0.75</b>	0.56	0.07	0.02	0.30	-0.00	0.23
Train Period	2016-2024	0.58	0.57	0.50	<b>0.54</b>	0.52	0.64	0.64	0.59	0.75	<b>0.72</b>	0.07	0.02	0.25	0.01	0.21

Freddie Mac		Test Period										Difference				
		No Offset					With Offset									
		2016-2017	2018-2019	2020-2021	2022-2023	2016-2024	2016-2017	2018-2019	2020-2021	2022-2023	2016-2024	2016-2017	2018-2019	2020-2021	2022-2023	2016-2024
Train Period	2016-2017	0.58	0.56	0.43	0.50	0.45	0.61	0.55	0.71	0.48	0.66	0.04	-0.01	0.28	-0.03	0.21
2018-2019	0.57	0.56	0.44	0.50	0.45	0.63	0.58	0.72	0.51	0.69	0.06	0.02	0.28	0.02	0.23	0.23
2020-2021	0.46	0.45	0.49	0.37	0.43	0.55	0.47	0.47	0.72	0.47	0.66	0.09	0.02	0.23	0.10	0.23
2022-2023	0.50	0.51	0.40	0.53	0.44	0.61	0.56	0.73	0.53	0.70	0.12	0.05	0.32	0.00	0.26	0.22
Train Period	2016-2024	0.52	0.51	0.47	0.50	0.47	0.61	0.54	0.73	0.51	0.69	0.09	0.02	0.26	0.02	0.22

Source Comparison		Test Period										Difference				
		No Offset					With Offset									
		2016-2017	2018-2019	2020-2021	2022-2023	2016-2024	2016-2017	2018-2019	2020-2021	2022-2023	2016-2024	2016-2017	2018-2019	2020-2021	2022-2023	2016-2024
Train Period	2016-2017	0.04	0.03	0.02	0.03	0.03	0.05	0.05	0.02	0.05	0.04	0.01	0.03	0.03	0.04	0.04
2018-2019	0.03	0.05	0.02	0.04	0.04	0.03	0.03	0.05	0.01	0.03	0.05	0.01	0.03	0.03	0.05	0.05
2020-2021	0.06	0.08	0.03	0.05	0.05	0.06	0.08	0.03	0.01	0.03	0.05	0.01	0.03	0.03	0.05	0.05
2022-2023	0.08	0.07	0.03	0.05	0.05	0.06	0.04	0.04	0.02	0.04	0.06	0.01	0.04	0.04	0.03	0.03
Train Period	2016-2024	0.06	0.06	0.02	0.05	0.05	0.06	0.04	0.02	0.04	0.06	0.01	0.02	0.04	0.03	0.03

modeled. This mortgage data, where the scoring attributes are either linearized or predefined discrete factor levels, eliminates the need for SGBRT’s nonlinear flexibility.

Overall, our results confirm that incorporating APC inputs into both panel logistic regression and stochastic gradient boosted regression trees significantly improves model performance during economic transitions, particularly the unprecedented conditions of the COVID-19 pandemic. This approach not only allowed for the inclusion of pandemic data in the modeling but actually improved the resulting scores compared to excluding that data.

## 5 Conclusion

The fragility of credit scores to economic changes has often been accepted as unavoidable. This paper demonstrates that credit scores can be designed to be resilient through economic volatility and, in fact, benefit from volatility. The keys to this resilience are 1) using panel data (one observation per account per month) rather than cross-sectional data (one observation per account for the full training window), and 2) integrating Age-Period-Cohort lifecycle and environment into the credit score estimation, whether traditional logistic regression or machine learning.

This approach was quite effective for being able to model data from the worst economic period of the COVID-19 pandemic. One particularly powerful aspect is that no macroeconomic variables were used. The pandemic is notorious for breaking the traditional correlations between macroeconomic factors and loan defaults due to the unprecedented government intervention and forbearance programs. The APC environment function can quantify what the environment was during any given time period without needing to explain what caused that environment in macroeconomic terms.

Although the current research focused on modeling defaults, the same methods can be applied to other key metrics, most commonly prepayment, attrition, and recovery rates. Although survival models are designed to model terminal events, the APC approach applies to any variable. For products such as credit cards, this can include modeling purchase, payment, utilization, and revolving balance rates.

The work here was performed on Fannie Mae and Freddie Mac mortgage default data. Similar analysis for auto, credit card, and personal loans internally at banks with proprietary data has shown even more dramatic benefits

because of the smaller, more dynamic portfolios being modeled.

The specific combination of APC and credit scoring is one of a family of algorithms that could be employed to achieve similar goals. Broadly speaking, combining discrete time survival models and machine learning models can overcome many shortcomings of today's credit scores.

## References

- [1] Erik Andersson. *Statistical Credit Rating with Survival Regression & Gradient Boosting*. PhD thesis, KTH Royal Institute of Technology, 2023.
- [2] Alberto Archetti and Lorenzo Noci. Fpboost: Fully parametric gradient boosting for survival analysis. *arXiv preprint*, 2024.
- [3] Yacine Aït-Sahalia, Dacheng Xiu, and Yilmaz Yilmaz. Covid-19 and credit risk: A long memory perspective. *Journal of Financial Economics*, 145(2):402–425, 2022.
- [4] Avinash Barnwal, Hyunsu Cho, and Toby Hocking. Survival regression with accelerated failure time model in xgboost. *Journal of Machine Learning Research*, 22:1–22, 2021.
- [5] Anthony Bellotti and Jonathan Crook. Survival analysis in credit risk management: A review study. *Journal of Credit Risk*, 20(3):45–78, 2024.
- [6] Sudipta Bose, Syed Shams, Muhammad Jahangir Ali, and Dessalegn Mihret. Covid-19 pandemic risk and probability of loan default. *Finance Research Letters*, 43, 2021.
- [7] Joseph L. Breeden. Validation of stress testing models. In George A. Christodoulakis and Stephen Satchell, editors, *The Analytics of Risk Model Validation*, pages 13–25. Academic Press, 2008.
- [8] Joseph L. Breeden. Macroeconomic adverse selection: How consumer demand drives credit quality. In *Credit Scoring and Credit Control XII Conference*, Edinburgh, August 2011.
- [9] Joseph L. Breeden. *Living with CECL: Mortgage Modeling Alternatives*. Prescient Models LLC, 2018.

- [10] Joseph L. Breeden. *Redesigning Credit Risk Modeling to Achieve Profit and Volatility Targets*. Bookbaby Publishing, USA, 2024.
- [11] Joseph L. Breeden, Anthony Bellotti, and Yevgeniya Leonova. Instabilities in cox proportional hazards models in credit risk. *Journal of Credit Risk*, 19(2):29–55, May 2023.
- [12] Joseph L. Breeden and Jose J. Canals-Cerdá. Consumer risk appetite, the credit cycle, and the housing bubble. *Journal of Credit Risk*, 14(2):1–30, 2018.
- [13] Joseph L. Breeden and Yevgeniya Leonova. Macroeconomic adverse selection in machine learning models of credit risk. *Engineering Proceedings*, 39(1), 2023.
- [14] Joseph L. Breeden and Yevgeniya Leonova. Stabilizing machine learning models with age-period-cohort inputs for scoring and stress testing. *Frontiers in Applied Mathematics and Statistics*, 9:1195810, 2023.
- [15] Paul S Calem, Matthew Cannon, and Leonard I Nakamura. Credit cycle and adverse selection effects in consumer credit markets-evidence from the heloc market. Technical report, Board of Governors of the Federal Reserve Bank, 2011.
- [16] Alberto Cavallo. Inflation with COVID consumption baskets. Working Paper 27352, National Bureau of Economic Research, 2020.
- [17] Tianqi Chen and Carlos Guestrin. A gradient boosting algorithm for survival analysis via dynamic programming. *arXiv preprint*, 2013.
- [18] D. R. Cox. Regression models and life-tables. *Journal of the Royal Statistical Society. Series B (Methodological)*, 34(2):187–220, 1972.
- [19] J. Davis, L. Devos, S. Reyners, and W. Schoutens. Gradient boosting for quantitative finance. 2019.
- [20] Deloitte Insights. Financial well-being index (FWBI): Tracking consumer financial health. Technical report, Deloitte Development LLC, 2024.

- [21] W. Erwin Diewert and Kevin J. Fox. Measuring real consumption and CPI bias under lockdown conditions. *Health Economics*, 31(S1):114–135, 2022.
- [22] Lore Dirick, Gerda Claeskens, and Bart Baesens. Time to default in credit scoring using survival analysis: A benchmark study. *Journal of the Operational Research Society*, 68(6):652–665, 2017.
- [23] Fannie Mae. Fannie mae single-family loan performance data, 2020.
- [24] Freddie Mac. Single family loan-level dataset, 2025.
- [25] Jerome H Friedman. Stochastic gradient boosting. *Computational statistics & data analysis*, 38(4):367–378, 2002.
- [26] Wenjiang Fu. *A Practical Guide to Age-Period-Cohort Analysis: The Identification Problem and Beyond*. Chapman and Hall/CRC, 2018.
- [27] T R Holford. The estimation of age, period and cohort effects for vital rates. *Biometrics*, 39(2):311–324, 1983.
- [28] Theodore Holford. *Encyclopedia of Statistics in Behavioral Science*, chapter Age-Period-Cohort Analysis. Wiley, 2005.
- [29] Kearney Consumer Institute. Consumer stress index: Understanding the gap between sentiment and spending. Technical report, Kearney, 2025.
- [30] Lei Li and Philip E. Strahan. Who supplies PPP loans (and does it matter)? banks, relationships, and the COVID crisis. *Journal of Financial and Quantitative Analysis*, 56(7):2411–2438, 2021.
- [31] Daniel J Sargent. A flexible approach to time-varying coefficients in the cox regression setting. *Lifetime Data Analysis*, 3(1):13, 1997.
- [32] Volker Schmid and Leonhard Held. Bayesian age-period-cohort modeling and prediction - bamp. *Journal of Statistical Software, Articles*, 21(8):1–15, 2007.
- [33] Maria Stepanova and Lyn C. Thomas. Dynamic survival models with time varying coefficients for credit risks. *Journal of Operational Risk*, 18(2):1–24, 2023.

- [34] Mark Thackham. *Survival Analysis Applications to Credit Risk Default Modelling*. PhD thesis, Macquarie University, 2022.
- [35] Terry M. Therneau and Patricia M. Grambsch. *Modeling Survival Data: Extending the Cox Model*. Springer-Verlag, New York, 2000.
- [36] Lyn C. Thomas and Jonathan Crook. Development and validation of credit-scoring models using panel data. Technical report, Cornell University, 2007.
- [37] Yufei Xia, Chuanyi Liu, Yinguo Li, and Na Liu. A dynamic credit scoring model based on survival gradient boosting decision tree approach. *Technological and Economic Development of Economy*, 27(1):96–119, 2021.
- [38] Y. Yang and K.C. Land. *Age-Period-Cohort Analysis*. Taylor and Francis, Boca Raton, 2014.