Federal Reserve Bank of Cleveland Working Paper Series

Macroeconomic Parameter Instability in Auto Loan Loss Models

Nicholas T. Fritsch and Edward S. Prescott

Working Paper No. 26-18

July 2026

Suggested citation: Fritsch, Nicholas T. and Edward S. Prescott. 2026. "Macroeconomic Parameter Instability in Auto Loan Loss Models." Working Paper No. 26-18. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202618.

Federal Reserve Bank of Cleveland Working Paper Series
ISSN: 2573-7953

Macroeconomic Parameter Instability in Auto Loan Loss Models

Nicholas T. Fritsch*
Federal Reserve Bank of Cleveland
Edward S. Prescott*
Federal Reserve Bank of Cleveland

July 8, 2026

Abstract

We estimate a discrete-time Markov transition model of auto loan performance over the 2000-2025 period using multinomial logistic regressions. Using rolling pseudo out-of-sample forecasts, we document persistent declines in the sensitivity of the probability of default to an unemployment shock in both the global financial crisis and Covid recessions. The estimated effect of a 1 percentage point increase in unemployment on default risk declined from 16 percent in 2006 to 3 percent post-2020. Two-year cumulative default forecasts over 2020-2021 using pre-pandemic parameters overstate actual defaults by 100 basis points (25 percent), with forecast errors largest in absolute terms for subprime borrowers and largest in relative terms for prime borrowers. The instability persists after controlling for forbearance usage and pandemic-period dummy variables, and is driven primarily by changes in macroeconomic relationships rather than borrower composition. These findings have implications for stress testing models and loss forecasting practices that rely on stable unemployment-default relationships.

Keywords: Auto loans, parameter instability, stress testing, default forecasting, Markov transition models

JEL Codes: G21, G28, C25, C53

*The views expressed in this paper are those of the authors and not necessarily those of the Federal Reserve Bank of Cleveland or the Federal Reserve System. This paper was presented at the Federal Reserve Bank of Cleveland. We would like to thank Rebel Cole, Ben Craig, Joe Haubrich, Niko Hauzenberger, Gary Koop, Lara Loewenstein, and Stuart McIntyre for helpful comments. Corresponding author: Nicholas.Fritsch@clev.frb.org.

1 Introduction

This paper estimates a discrete-time Markov transition model of auto loan performance using multinomial logistic regressions. The model is estimated over the 2000-2025 period, which includes the global financial crisis (GFC), and the Covid shock. The first purpose of this paper is to evaluate the stability of estimated coefficients on macroeconomic shocks like unemployment. These coefficients matter for the stress testing models that banking regulatory agencies run to evaluate bank capital and they matter for loss forecasting by auto lenders. Using rolling regressions we demonstrate that these coefficients change significantly in response to the large aggregate shocks of the GFC and the Covid periods and the effect on these estimates is persistent.

Figure 1 shows the instability of the coefficient on our unemployment variable for rolling regression estimates of the relative risks of a current loans going into default. The effects are particularly evident for the Covid period, when despite large spikes in unemployment, auto lending default rates declined. While not necessarily surprising due to the large government interventions in this period, these effects are also evident during the 2007-2009 recession as well. These findings are quantitatively significant. For example, our model produces forecast errors of up to 100 basis points in cumulative default rates over a two-year horizon during the Covid period, representing approximately a 25 percent overstatement relative to actual outcomes.

The second purpose of this paper is to apply discrete-time Markov transition methods to modeling auto default. This technique has been used in the mortgage lending literature, but less so in the auto lending literature. In auto lending, many loans become delinquent before going into default, but a significant fraction either exit for non-default reasons such as prepayment or return to being current after becoming delinquent. Both prepayment and default should be jointly accounted for to avoid potentially large errors in estimating transition probabilities (Deng et al., 2000). The Markov transition framework jointly estimates the transition probabilities between all the states and uses the information in these transition estimates to forecast defaults several periods ahead.

The final purpose of this paper is to add to the auto lending literature by demonstrating the importance of different lending channels for default, that is, whether a loan is made by a credit union, bank, or auto finance company. We document significant segmentation in auto lending by lender type with subprime lending concentrated in the auto finance sector. Furthermore, even when observables are accounted for, loan performance differs by lender type. For example, subprime loans originated by auto finance lenders are significantly riskier than subprime loans originated by other channels. The safest loans are those originated by credit unions and banks. We are aware of only two other papers that document heterogeneity in auto loan performance between banks, credit unions, and auto finance lenders.[1]

2 Background on the Auto Loan Market

The auto loan market is the second largest consumer credit market in the United States. In the fourth quarter of 2025 there was $1.56 trillion of outstanding auto loans.[2]

Most auto loans are made through car dealers. A person buying a car goes to a dealer, and after the person decides to purchase a car, the dealer will present the buyer with some auto loan options. Dealers work with a variety of lenders. However, there are significant differences in auto loans by type of lender. The left panel of Figure 2 shows the percent of loans originated by lender types in our data, including auto finance companies, banks, and credit unions. Since 2000, banks' market share of the auto loan market has declined, while credit unions' market share has increased steadily. Prior work has documented that auto finance companies specializing in lending to the subprime segment of the market played an increasing role in financing that borrower segment during the pandemic (Canals-Cerda and Lee, 2021). The right panel of Figure 2 shows that auto finance companies originate the largest share of subprime auto loans.

[1] A few papers in the literature have focused on lender effects, but either did not identify the lender as a bank, credit union, or auto finance lender, or were focused on loan pricing rather than loan performance.

[2] Z.1 Financial Accounts of the United States.

Not only has the composition of lenders changed over time, but so has the distribution of auto loan origination credit risk. Figure 3 plots the Equifax Risk Score (henceforth "risk score") in our data from 2000-2025. Over that period, the aggregate share of loans in higher risk score groups has increased. Roughly 20 percent of loans were above a risk score of 760 in 2000, but this increased to roughly 40 percent of loans by 2025.

Historically, auto performance has followed the economy. Figure 4 shows the actual transition rates over 2000-2025. This period covers three recessions. The mild 2001 recession, the deep recession of 2007-2009, and the short Covid recession in 2020.[3] Unlike in the other two recessions, auto defaults did not increase in 2020 and 2021; instead, default rates decreased. The decoupling of unemployment and auto loan risk outcomes in 2020 has been attributed to large government transfers and lender forbearance policies that appeared to mitigate credit risks (for example, Bakshi and Rose (2021); Dettling and Lambie-Hanson (2021)). Forbearance usage, in particular, rose to historically high levels in 2020 (Figure 5). These policies appeared to have distorted the traditional relationship between auto loan performance and unemployment. Some of this was due to forbearance, but as we will discuss, most of the better performance was not.

3 Literature

The literature on auto loan default probabilities is small, in part due to the lack of loan-level auto data that have been historically available to researchers. The earliest papers include Agarwal and Ambrose (2007), Agarwal (2008), and Heitfield and Sabarwal (2004), who model auto default and prepayment as a function of loan age, loan and borrower characteristics, and macroeconomic variables.[4]

[3]Transition rates for the first few periods in our sample are distorted due to censoring in the first few years of our dataset.

[4]Motivated by the recent increase in the length of auto loan terms, several recent studies find that loan age is important in explaining default risk, with loans of greater than five years having worse outcomes (Katcher et al. (2024), Guo et al. (2022); Wu et al. (2018)). We do not analyze this loan characteristic due to data limitations in our dataset.

Broadly, the literature on loan loss forecasting and stress testing find that unemployment and other macroeconomic variables are important for modeling consumer loan default.[5] Most auto lending models include macroeconomic shocks. For example, Agarwal (2008) and Heitfield and Sabarwal (2004) found that default rates are sensitive to unemployment shocks, but prepayment rates are less so. Heitfield and Sabarwal (2004) also found that auto loan prepayment rates are insensitive to market interest rates. The Federal Reserve auto lending stress test includes unemployment and house price changes in its model.[6]

One challenge with modeling the effects of macroeconomic variables is the potential instability of estimates. Gross and Población (2019) and Guerrieri and Welch (2012) document that macroeconomic variables can lead to large amounts of uncertainty in reduced form statistical models of credit risk and bank performance, so should be used with caution. The study of auto lending stress tests by Wu et al. (2018) shows significant differences in estimates of coefficients on macro variables in pre-GFC and post-GFC samples, though they use a difference dataset for each period. They speculate that this is because they have only one recession in their sample. Unlike Wu et al. (2018) we use a single dataset and show that this result is not due to different data sources. Furthermore, we use rolling regressions that show how the estimate changes over time and why it changes. Finally, since our analysis covers two recessions, it shows that this instability is not unique to one recession.[7]

Much of the auto lending literature estimates discrete-time models that focus only on the exit states, such as conditional probability of default (PD) or competing risks models (An et al. (2020); Agarwal (2008); Heitfield and Sabarwal (2004); Guo et al. (2022); Wu et al. (2018)). The literature that applies Markov transition matrices to loan-level data typically does so on mortgage data (Smith and Lawrence (1995) and Smith et al. (1996), Chen et al. (2020)). One exception in the auto lending literature is Smith and Jin (2007). While that

[5] See, for example, Covas et al. (2014) and Hirtle et al. (2016).

[6] See https://www.federalreserve.gov/supervisionreg/files/credit-risk-models.pdf.

[7] The earliest part of the sample also covers the 2000-2001 recession, but we start our rolling regressions in 2006 so that we have at least one recesssion in the first rolling estimation. This produces a single set of coefficients for the period covering 2000-2006, so we do not examine potential instability after the 2000-2001 recession.

paper estimates Markov transition matrices, it does so on auto lease data and does not examine the effect of macroeconomic variables.

The last strand of the literature related to our work covers risk from different lender channels in auto lending. An et al. (2020) find that the lender channel is important in explaining default probabilities, with credit unions and banks facing relatively less risk of default when compared to auto finance companies. While nearly all studies establish a link between borrower credit quality and default probability, Heitfield and Sabarwal (2004) discuss the heterogeneous risks that exist for lenders who target different segments of the subprime market. Katcher et al. (2024) document differential prepayment rates and relationships between maturity and interest rates across lender types against a backdrop of increasing auto loan maturities over time. Our analysis builds on this analysis by further analyzing the effect of the lender channel.

4 Data

We use loan-level data on automobile loans from the FRBNY Consumer Credit Panel / Equifax Data (CCP) from June 2000-June 2025. This panel is an anonymized 5 percent national sample of all consumers who have a credit history reported to Equifax. The data contain auto loan and borrower characteristics over the life of the loan.

The data are semi-annual prior to 2017, quarterly over 2017-2019, and monthly beginning in 2020. We use semi-annual samples of the data for our analysis to be consistent throughout our sample period. The semi-annual samples in the CCP are as of June 1 and December 1 in each year. Furthermore, prior to 2015, the dataset only reports three broad types of lenders: auto finance companies, banks, and credit unions, so we use this classification of lender type.

We take several steps to clean the data prior to model estimation. First, the consumers the CCP follows are called primary borrowers. Some households include other people who also borrow. These other borrowers are not necessarily followed over time, so we drop them from our sample. Second, we drop loans from US territories, because these areas have few observations in the data. Third, we remove observations with missing or zero values for borrower age, Equifax Risk Score, or outstanding balances. Fourth, we remove leases, loans observed for greater than 9 years, loans where the borrower died during the life of the loan, loans with duplicated identifiers in a given time period, and a small number of loans with original balances of less than \$5,000 or more than \$100,000.[8]

We categorize a loan as being in one of four possible states. A loan is *current* if the payment is less than 30 days overdue. A loan is *delinquent* if the loan payment is 30 - 89 days late. A loan is in *default* if it is at least 90 days delinquent, repossessed, bankrupt, or charged off. Finally, a loan is a *non-default exit* if it exits and did not default. Non-default exit may include loans in which the balance is fully paid at term, is prepaid or refinanced, the borrower died, or the loan was sold by the lender. The reason for a non-default exit is usually unobservable in our data.[9] Furthermore, it is possible that some of these non-default exits are really defaults due to a loan going delinquent and exiting between periods of observation. To account for this possibility, we follow the procedure used by the Federal Reserve stress test, which is to make assumptions based on the size and status of the loan in the previous period along with the borrower's credit score.[10] Once a loan has exited the sample for default or non-default reasons, it exits our dataset.

We use borrower characteristics, loan features, and state-level macroeconomic variables in our estimation. The borrower characteristics are age, Equifax Risk Score, and whether the borrower has a mortgage. The loan features are the age of the loan, the outstanding loan balance, whether the loan is held jointly with someone else, and the type of lender. Outstanding loan balances are deflated by the consumer price index (CPI) using June 1999 as the base period. Other than the lender type, these variables are similar to the variables

[8]For example, the largest 0.001 percent of loans in the data carried balances greater than \$90,000.

[9]The CCP does not always report if a borrower dies. When it is reported, we drop the loan.

[10]For the details of the Federal Reserve's process, see documentation from Board of Governors of the Federal Reserve System https://www.federalreserve.gov/supervisionreg/files/credit-risk-models.pdf.

used in the literature (Agarwal (2008); An et al. (2020); Guo et al. (2022); Heitfield and Sabarwal (2004); Wu et al. (2018)). The loan and borrower age, mortgage status, and loan balance for each loan are allowed to vary over time for each loan in our model. We use the loan origination risk score of the borrower because movements in risk score may correlate with changes in other model variables, such as macroeconomic outcomes. We categorize the origination risk score into groups that broadly correspond to credit risk groupings typically used by lenders, spanning from super prime (an Equifax Risk Score greater than 760) to deep subprime (an Equifax Risk Score less than 560).[11] All other loan and borrower characteristics in the model are constant over the life of the loan.

The state-level macroeconomic variables we use are seasonally adjusted unemployment data from the Bureau of Labor Statistics (BLS) and the all-transactions house price index (HPI) from the Federal Housing Finance Agency. We follow Wu et al. (2018) and measure these variables in terms of year-over-year changes. The Federal Reserve's stress test also uses year-over-year changes in these two variables in its stress testing model, but their model also includes the level of unemployment as well as interactions of these variables with credit scores.[12] In our model, we find a negative coefficient on the level of unemployment associated with transitions to default, so we drop it from our model. Furthermore, given the number of other variables in our model, we do not interact unemployment change with credit score groups. We also exclude state fixed effects from our model because we find that these impair the accuracy of model predictions, and are likely highly correlated with the state-level macroeconomic variables. Finally, we tried several other potentially relevant macroeconomic variables such as Treasury rates, the prime rate, and used car price indices, but these were not statistically significant or meaningful in our estimation results, so we leave them out.

After cleaning and data construction, we have 296,388 unique loans and 1,668,736 total

[11] Because the Equifax risk score is based on an internal model used by Equifax, we consider the possibility that the groupings we use may not match the credit score ranges commonly used. For robustness, we also categorize borrowers based on the origination quintiles of their Equifax Risk Score. Our results do not change using either approach.

[12] See https://www.federalreserve.gov/supervisionreg/files/credit-risk-models.pdf.

observations. Table 1 describes the variables used in our empirical model. Table 2 provides summary statistics for the transition rates and variables used in the regression analysis over the full sample from 2000-2025.

5 Empirical Framework

We use a discrete-time Markov transition framework where conditional probabilities of transition are estimated using a multinomial logit model. As discussed above, this approach has been used in the mortage literature, but less so in the auto lending literature.

The benefit of the transition matrix framework is that it takes into account movement through the delinquency state and into the non-default exit state. Default modeling is often used to construct a measure of cumulative default over a longer horizon in order to assess the need for additional loss-absorbing capital. Multi-period-ahead forecasting using the transition matrix approach takes into account the probabilistic movements between states when constructing a cumulative default probability.

The Markov approach uses a recursive structure for forecasting the likelihood of transitioning between states. For each set of possible borrowers, lenders, and macroeconomic shocks let $\boldsymbol{x_{it}}$ be a vector of characteristics for loan i and let $\boldsymbol{z_t}$ be a vector of macroeconomic shocks, at time t. Let $\mathbf{y_t}$ be the vector of possible states for a loan. Conditioning the state transitions on $\boldsymbol{x}$ and $\boldsymbol{z}$, the Markov property assumes that

$$\begin{aligned} prob(\boldsymbol{y_{t+1}} = k|\boldsymbol{y_t} = j, \boldsymbol{y_{t-1}} = j_{t-1}, ..., \boldsymbol{y_{t-k}} = j_{t-k}; \boldsymbol{x_{it}}, \boldsymbol{z_t}) \\ = prob(\boldsymbol{y_{t+1}} = k|\boldsymbol{y_t} = j; \boldsymbol{x_{it}}, \boldsymbol{z_t}) = p_{jk}(\boldsymbol{x_{it}}, \boldsymbol{z_t}). \end{aligned}$$

We let $\boldsymbol{P}(\boldsymbol{x_{it}}, \boldsymbol{z_t})$ be the matrix of transition probabilities as a function of loan and borrower characteristics and macroeconomic shocks at time t.

In our analysis, there are four states. A loan can be current (c), delinquent (dq), default (df), or a non-default exit (e). The default and non-default exit states are absorbing states.

The transition matrix $\boldsymbol{P}(\boldsymbol{x_{it}}, \boldsymbol{z_t})$ contains elements $p_{jk}(\boldsymbol{x_{it}}, \boldsymbol{z_t})$ that describe each of the possible transition probabilities at time t,

$$\mathbf{P}(\mathbf{x_{it}}, \boldsymbol{z_t}) = \begin{pmatrix} p_{c,c}(\boldsymbol{x_{it}}, \boldsymbol{z_t}) & p_{c,dq}(\boldsymbol{x_{it}}, \boldsymbol{z_t}) & p_{c,df}(\boldsymbol{x_{it}}, \boldsymbol{z_t}) & p_{c,e}(\boldsymbol{x_{it}}, \boldsymbol{z_t}) \\ p_{dq,c}(\boldsymbol{x_{it}}, \boldsymbol{z_t}) & p_{dq,dq}(\boldsymbol{x_{it}}, \boldsymbol{z_t}) & p_{dq,df}(\boldsymbol{x_{it}}, \boldsymbol{z_t}) & p_{dq,e}(\boldsymbol{x_{it}}, \boldsymbol{z_t}) \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{1}$$

At this stage, the problem is to estimate the transition matrices $\mathbf{P}(\mathbf{x}_{\boldsymbol{it}}, \mathbf{z}_{\boldsymbol{t}})$. Given the need to estimate probabilities across states and the fact that the data represent semiannual observations, we estimate conditional transition probabilities using a discrete time multinomial logit model. The model assumes that the conditional log odds of transitioning to state k depend linearly on a vector of loan and borrower characteristics $\boldsymbol{x_{it}}$ for loan i and a vector of macroeconomic variables $\boldsymbol{z_t}$ with corresponding coefficient vectors $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$. Let $p_{i,jk}(\boldsymbol{x_{it}}, \boldsymbol{z_t})$ be the probability that loan i will transition from state j to state k at time t given $\boldsymbol{x_{it}}$ and $\boldsymbol{z_t}$. The model estimates

$$ln(\frac{p_{i,jk}(\boldsymbol{x_{it}}, \boldsymbol{z_t})}{p_{i,jj}(\boldsymbol{x_{it}}, \boldsymbol{z_t})}) = \boldsymbol{x_{it}\beta_{jk}} + \boldsymbol{z_t\gamma_{jk}} \qquad \text{for } k \neq j, \tag{2}$$

which gives

$$\frac{p_{i,jk}(\boldsymbol{x_{it}}, \boldsymbol{z_t})}{p_{i,jj}(\boldsymbol{x_{it}}, \boldsymbol{z_t})} = exp(\boldsymbol{x_{it}\beta_{jk}})exp(\boldsymbol{z_t\gamma_{jk}}) \qquad \text{for } k \neq j$$

Exponentiation of the dependent variable is the probability of transition to state k relative to the probability of staying in the baseline state j. This *relative risk* (RR) is an intuitive way to interpret the coefficient estimates of the multinomial logit. A coefficient greater than 1 means that the relative probability of transitioning to a state increased more than the probability of staying in the same state, holding other covariates fixed. We define it as

$$RR_{i,k}(\boldsymbol{x_{it}}, \boldsymbol{z_t}) = \frac{p_{i,jk}(\boldsymbol{x_{it}}, \boldsymbol{z_t})}{p_{i,jj}(\boldsymbol{x_{it}}, \boldsymbol{z_t})} \qquad \text{for } k \neq j.$$

The effect of a 1-unit change in x_{it} on RR is

$$\frac{RR_{i,k}(x_{it}+1, \boldsymbol{z_t})}{RR_{i,k}(x_{it}, \boldsymbol{z_t})} = \frac{exp(\beta_{ik}(x_{it}+1))}{exp(\beta_{ik}x_{it})} = exp(\beta_{ik})$$

with a similar equation for the effect of a 1-unit change in z.

Finally, solving for the conditional transition probabilities that enter into $\boldsymbol{P}(t)$,

$$p_{i,jk}(\boldsymbol{x_{it}}, \boldsymbol{z_t}) = \frac{exp(\boldsymbol{x_{it}\beta_{jk}})exp(\boldsymbol{z_t\gamma_{jk}})}{1 + \sum_{k \neq j} exp(\boldsymbol{x_{i,t-1}\beta_{jk}})exp(\boldsymbol{z_t\gamma_{jk}})} \tag{3}$$

$$p_{i,jj}(\boldsymbol{x_{it}}, \boldsymbol{z_t}) = \frac{1}{1 + \sum_{k \neq j} exp(\boldsymbol{x_{it}\beta_{jk}})exp(\boldsymbol{z_t\gamma_{jk}})} \tag{4}$$

The multinomial logit model jointly estimates parameters of each transition equation such that the estimated probabilities are constrained to be between 0 and 1 and sum to 1. Because we are modeling the entire state space, we estimate a separate multinomial logistic regression for the two starting states, current and delinquent. The estimated probabilities correspond to the first two rows of (1). In all estimations, we cluster standard errors at the loan level.

5.1 Multi-Period-Ahead Forecasts of Default

Estimating a cumulative probability of default or non-default exit over multiple periods is commonly used by lenders for loss forecasting and by regulators for stress testing. This exercise is straightforward in the Markov transition matrix framework, but requires assumptions on the data given that the probabilities are conditional on changing loan and macroeconomic variables. One method is to assume a path of macroeconomic variables while holding other non-deterministic variables constant. In this setting, lenders and regulators can assess the sensitivity of existing loan portfolios to macroeconomic shocks.

To estimate a cumulative probability of default and non-default exit, we need the initial distribution of states that loans are in, $\pi_{\mathbf{0}}$; the initial distribution of loan and borrower characteristics, $(\mathbf{x_{i0}}, \mathbf{z_0})$; how the deterministic characteristics change over time, e.g., loans age deterministically, which we call $\tilde{\mathbf{x}}_{\mathbf{il}}$ to reflect that only certain characteristics change over time; the time path of the macro shocks $\tilde{\mathbf{z}}_{\mathbf{l}}$; and the transition matrix P. The h-period-ahead cumulative forecast of the distribution of states, $\pi_{\mathbf{0+h}}$ is calculated by

$$\pi_{\mathbf{0+h}} = \pi_{\mathbf{0}}\mathbf{P}(\mathbf{x_{i0}}, \mathbf{z_0})\mathbf{P}(\tilde{\mathbf{x}}_{\mathbf{i1}}, \tilde{\mathbf{z}}_{\mathbf{1}}) \cdots \mathbf{P}(\tilde{\mathbf{x}}_{\mathbf{ih}}, \tilde{\mathbf{z}}_{\mathbf{h}}) \tag{5}$$

6 Model Estimation

For loan age, we use a quadratic specification of the hazard as a function of age. As Lancaster (1990) demonstrates, a misspecified hazard function can fundamentally alter conclusions derived from an estimated model. We compared the quadratic specification with several other functional forms and found that it had the best fit. See Appendix A.1 for the details.

Using our chosen hazard specification, we first estimate the full model specification from 2000-2019 in order to understand relationships between model covariates and transitions prior to potential distortions from the pandemic period. Tables 3 and 4 show the results of our estimation. In each table the coefficients represent the marginal effects on the relative risk of transition to a particular state compared with staying in the current state. For the categorical variables that we include in the model that have more than two levels, credit unions and origination risk scores greater than 760 are the reference groups for lender type and origination risk score, respectively.

6.1 Transitions from the Current State

We first focus on the estimates of the transition from the current state. The loan and borrower characteristics included in the model are almost all statistically significant at the 1 percent level. Borrowers who are younger, in lower risk score groups, non-mortgage holders, further from origination, and do not have a joint account have higher relative risks of transition into delinquency or default relative to remaining current. The impacts of low risk scores are particularly pronounced. For example, the odds of a deep subprime borrower transitioning from current to delinquency (default) relative to a transition to current for borrowers with a risk score of at least 760 increase by a factor of 7000 percent (2200 percent). The lender channel also shows meaningful effects, particularly for loans originated by auto finance companies. The odds of an auto finance loan transitioning to delinquency (default) relative to staying current and originating from a credit union are 70 percent (55 percent) higher.

Borrowers with the highest origination risk scores, mortgage holders, and borrowers with lower balances have relatively higher odds of transitioning to non-default exit. The absolute magnitude of estimated effects for risk score and lender type on the relative risk of transitioning to the non-default exit state is smaller than that those for transitioning to default. The odds of a deep subprime (<560) borrower transitioning to exit relative to a transition to current for borrowers with a risk score of at least 760 decrease by 17 percent. The relative risks of exit for loans that originate from auto finance lenders are 3 percent lower relative to staying current and originating from a credit union.

Unlike the coefficients on risk score and lender type, the coefficients on macroeconomic variables are not always statistically significant. A 1 percentage point increase in the year-over-year change in the unemployment rate is associated with a 9 percent increase in the relative risk of transitioning to default, but the other states do not show statistical significance. A 1 percentage point increase in the annual HPI slightly decreases the relative risk of transitioning to delinquency, and slightly increases the relative risk of transitioning to non-default exit. While these effects seem small compared to credit score or lender effects, large macroeconomic shocks during economic downturns can have larger effects on the relative risks. Indeed, this is the primary role for these variable in stress testing models that condition on tail macroeconomic outcomes.

6.2 Transitions from the Delinquent State

For transitions from the delinquent state, estimated relative risk coefficients on the borrower and loan characteristics are again generally statistically significant and intuitive. The patterns in the current state model are similar to those in the delinquent state model, but borrower age no longer matters statistically or in a meaningful way. Also, unlike starting in the current state, an increase in the loan age results in a decrease in the relative risk of transitioning into default for borrowers in the delinquent state. A loan originated from auto finance lenders increases the relative risks of transitioning to delinquency and default by 9 and 10 percent, respectively, relative to transitions into current by borrowers from credit unions. The relative risks of transitions into delinquency and default for loans originated with a risk score of less than 560 increase by 284 and 254 percent, respectively, relative to transitions into current by borrowers with an origination credit score higher than 760. For transitions into non-default exit, the results reverse such that the relative risks of exit are decreased for loans originated from auto finance companies and for borrowers with lower credit risk scores.

The estimated effects of macroeconomic variables are not always significant or intuitive. For the year-over-year change in unemployment, the effect is only statistically significant for transitions to the exit state, but the effect suggests that an increase in unemployment increases the relative risks of non-default exit. As discussed earlier, it is possible that our data-cleaning procedure for identifying defaults from transitions to non-default exit misses some defaults. If that is the case, this measurement error could account for this finding. The estimated effects of a 1 percentage point increase in the annual HPI is statistically significant for the exit and default states, and suggests a slight increase in the relative risk of transition to exit and a slight decrease in the relative risk of transition to default.

7 Rolling Regressions, Predicted Transition Probabilities, and Coefficient Instability

We now run rolling regressions to estimate period-by-period aggregate transition probabilities and forecasted losses and then compare them to actual outcomes. We also examine how estimated coefficients and forecasted losses change as new data arrive and the model is re-estimated with the new data.

We use the period from 2000-2006 as the initial in-sample period. In our rolling predictions, we predict time $t+1$ outcomes on the parameter estimates estimated on all data up to time t. We then re-estimate the model parameters through all data as of $t+1$ and predict outcomes at $t+2$, and so on. This process continues until the last data point in our sample, ending in 2025.

Figure 6 plots the pseudo out-of-sample one-period-ahead forecasts (henceforth out-of-sample) compared with actual outcomes. The rows of the plots correspond to the first two rows of equation (1), summed over all loans at each time t. The red line in the first row plots the transition probability for current loans, given the characteristics of current loans at time t, and the red line in the second row plots the transition probabliity for delinquent loans, given the characteristics of delinquent loans at time t. Red shaded areas represent 95 percent confidence intervals around each prediction. The solid black line is the fraction of current or delinquent loans that actually transitioned to each new state from time t to time $t+1$. The dashed vertical line is December 2006, which is where the out-of-sample forecasts begin. (Forecasts in the plots prior to December 2006 are in-sample forecasts.)

The multinomial logit generally captures the patterns in the actual transition frequencies during the in-sample period and the out-of-sample rolling forecasts. After 2006, the model slightly underpredicts the transition from current to current in several periods, while it overpredicts the transitions from current to non-default exit. Similarly, the model slighly underpredicts transitions from delinquent to default, while it overpredicts from delinquent to exit.

The most notable forecast errors occur in the first semiannual period of 2020 for transitions from current to default and delinquent to exit. In this period, actual default rates declined, while the model predicted sharp increases in the transition probabilities. The next largest forecast errors occur during the GFC. Unlike during the GFC, the default and exit outcomes generally move in the same direction as the model predictions despite the model overpredicting the outcomes.

Forecast errors can arise when conditional relationships between the transitions and covariates fundamentally change. Given the size and scope of the federal government's response to Covid, this is a period where one would expect changes in these relationships. Our rolling coefficients allow us to examine how large these effects are on the estimated relationships. We focus on the coefficients in the current to default transition because current loans are a much larger share of loans than delinquent loans and for one-period-ahead forecasts it will have the largest impact on defaults.

We first examine the rolling standard deviation of the semiannual change in standardized rolling relative risk coefficients to look for evidence of coefficient volatility. Figure 7 shows the rolling standard deviations for transitions from current to default. With the exception of the model intercept, all of the loan and borrower covariates typically have a semiannual change of less than 1 standard deviation in 2020. The highest volatility in the first seminannual period of 2020 is in the seasonal fixed effect; presumably this rolling seasonal effect is picking up some of the pandemic-related changes during this period. The next highest volatility is in the unemployment coefficient, which changes by 1.3 standard deviations in June 2020 and the HPI change coefficient is similar. The volatility of both macro coefficients is usually small except during the GFC and Covid shocks. Figure 8 shows the rolling standard deviations for transitions from delinquent to exit. The coefficient instability for unemployment (nearly 3 standard deviations) was larger than for HPI.

While the coefficients for unemployment and HPI show volatility spikes in June 2020, we focus on unemployment for two reasons. First, unemployment experienced a much larger shock in 2020 (Figure 9) compared to the relatively stable HPI (Figure 10), which did not increase sharply until 2021. Second, the volatility of the unemployment coefficient is larger than that of the HPI coefficient when looking across both the GFC and Covid recessions, suggesting a more substantial change in default risk in response to unemployment shocks. This suggests unemployment as the primary driver of forecast errors during the two recessions in the rolling results.

Figure 11 reports the estimates of the unemployment coefficients through time in the rolling regressions. There is a large decline in this coefficient for the current-exit and current-default transitions during the GFC, and another sharp decline across four of the transition probabilities in 2020. (This drop explains the increase in volatility of the unemployment coefficient in Figures 7 and 8.) Given the large forecast errors associated with transitions from current-default in those two periods, we focus on the coefficient associated with that transition. At the beginning of the rolling estimation in December 2006, the marginal effect of a 1 percentage point increase in unemployment was associated with an increase in the relative risks of roughly 16 percent. After the GFC and before the pandemic, this effect drops to slightly less than 10 percent, which then declines again to about 3 percent in the post-pandemic period. Furthermore, the level of the effect after each decline is persistent, especially in the post-2020 period. After 2020, the coefficient stays low in the rolling regression through the end of our sample in 2025.

It is instructive to understand how these coefficients may have impacted the riskiest segments of the borrower population. We compute the marginal effects for a representative subprime borrower across a grid of values for annual unemployment changes to further illustrate the impact of an unemployment shock. We define a representative subprime borrower where the loan is originated by a bank and has an origination credit score of less than 560 while all other variables are held at their respective mean value across all borrowers. Figure 12 plots the increase in probability of transition from current to default for given year-over-year changes in US state unemployment. The x-axis represents a range of potential year-over-year changes in US state unemployment, the y-axis represents the marginal effect of a given change on the probability of transition from current to default.

In 2006, positive changes in unemployment of 10 percentage points (roughly the magnitude of the average unemployment change during the GFC) were associated with roughly a 1 percent increase in the probability of default. The size of the effect continues to increase as unemployment increases. Comparing this to the marginal effects generated across the same grid of annual unemployment increases in 2019 and 2021, it is only in 2006 that unemployment has a meaningful effect on default probabilities at relatively large shocks to the level of unemployment. There is a very small positive effect in 2019 between unemployment changes between 10 to 20 percent, but the effect is small relative to 2006. By 2021, the effect is effectively flat across the potential changes in annual unemployment.

These figures provide evidence of instability, namely, a decline in the relationship between unemployment and transitions to default after a large shock. The drop after 2020 should not be surprising. During Covid, there were two significant changes in the auto lending environment. First, there was a large spike in the use of forbearance by auto lenders. Second, the federal government made enormous transfers to households to offset the macroeconomic effects of Covid. These actions prevented many defaults, breaking down the traditional relationship between unemployment and transitions to default. Nevertheless, as we will see, these two factors do not change the underlying finding about instability.

7.1 Forbearance

One concern with our errors in forecasting default in 2020 is that they may be driven by borrowers who entered into forbearance. If characteristics of the borrower or loan are correlated with forbearance usage, or if forbearance usage occurred disproportionately in US states with larger spikes in unemployment, then some portion of our 2020 forecast errors could be a result of forbearance. As previously shown, forbearance increased to historically high levels in the same period of our default forecast errors.

To assess the degree to which loans that entered into forbearance change the forecast results, we compare the out-of-sample predicted probabilities between the full set of loans in our data and the set of loans that never enter into forbearance. Figure 13 shows the results. The red line indicates the average predictions from the rolling regressions on the full sample of loans; the blue line shows the average predictions for non-forbearance loans. The two forecasts almost perfectly overlap, including the period of the large current-default and delinquent-exit forecast errors in 2020. This suggests that forbearance does not explain our findings.

7.2 Controlling for 2020 Interventions

As we discussed, federal government interventions were enormous during Covid and far exceeded any other peacetime period in US history, which should alter the relationship between unemployment increases and defaults. To the extent that the pandemic represents an idiosyncratic period not representative of other recessionary periods, a simple way to control for this period is to include a dummy variable in the model for the pandemic period. We re-estimate (2) for the current and delinquent starting states, but now add a dummy variable equal to 1 during the semiannual periods from 2020 through 2021. If the dummy variable absorbs pandemic-specific factors that temporarily altered the unemployment-transition relationship, then the unemployment coefficients should better reflect the stable structural relationship that holds in normal periods.

Figure 14 shows the rolling coefficient estimates on the unemployment variable in the model that includes dummy variables for the pandemic period. We still find evidence of a declining relationship between unemployment and transitions to default. Relative to our prior results, the unemployment coefficient for the current-to-default transition declines more slowly over the 2020-2021 period and settles at a slightly higher level, but qualitatively the effect remains. We take this as evidence that the effects of the pandemic intervention are persistent on auto lending default performance, so a pandemic period dummy does not fully account for these effects.

8 Cumulative Forecasts

Lenders and regulators often construct cumulative forecasts of payoffs and default over time in order to project future revenues and the amount of loss-absorbing funds needed over a longer horizon. A two-year cumulative default rate is typically used in stress testing applications. For these types of exercises, our Markov structure is valuable. We assess the impact of the structural break in the relationship between unemployment and transitions to absorbing states on two-year cumulative forecasts. As we showed in Figure 6, forecasted transitions from current to delinquent were much higher than actual transitions. In this section, we break down the effect on forecasts from the structural break in the unemployment coefficient versus other coefficients.

In the predictions, we assume that the future path of macroeconomic variables is known, so that we are able to directly compare the model predictions to actual outcomes. For this exercise, additional assumptions are needed around the future path of the remaining model covariates. For borrower and loan characteristics that are deterministic (such as loan age and borrower age) we allow them to vary as time progresses in the forecast period, but non-deterministic variables that vary over time (such as outstanding loan balance and mortgage status) are held constant as of the latest observed data. Using these assumptions on the path of future data, (3) can be used to construct the time-varying predicted transition matrices and (5) can be used to construct a cumulative forecast based on those matrices.

We decompose the impact of the changing relationships between unemployment and transitions to default and exit by varying the parameters used to produce the cumulative forecasts. We use three sets of model parameters that we refer to as pre-date, post-date, and hybrid. The pre-date model coefficients are estimated as of December 2019. The post-date model uses the estimated coefficients as of June 2020. The hybrid model uses all of the post-date model parameters except for unemployment, and instead uses the unemployment coefficient as of December 2019. A comparison of these cumulative forecasts allows us to estimate how much of the variation in cumulative forecasts comes from revised relationships between unemployment versus the other model parameters.

Figure 15 shows the results of the default (left panel) and exit (right panel) forecasts as they accumulate from the first forecast period of June 2020 through December 2021, and Figure 16 plots the differences in the forecasts relative to actual outcomes at the end of the two-year period. The default predictions for the pre-date model simulates what was possible to forecast as of December 2019 if the path of unemployment and the HPI were known. The divergence between the pre-date model and actual outcomes by the end of the two-year cumulative forecast is roughly 100 basis points. As shown by Figure 16, this difference is economically significant and represents roughly a 25 percent overstatement in the predicted cumulative default relative to actual outcomes.

The predictions from the post-date model still produce a cumulative forecast higher than the actual outcome, but the differential is reduced by roughly two-thirds. Differences between the pre-date and post-date models point to the changes in coefficients as explaining much of the difference. The hybrid model breaks down the effect of changes in different coefficients. The hybrid model predictions are roughly halfway between the pre-date and post-date model predictions, which suggests that the change in the unemployment coefficient accounts for half of the difference in predictions and the other coefficients account for the other half.

8.1 Borrower Risk Score and Lender Channels

Our model controls for borrower risk score and origination lender type, so we use that to further examine how the cumulative forecasts differ across these sources of loan heterogeneity. Results along these dimensions will be correlated because banks and credit unions generally originate loans to higher quality borrowers relative to auto finance companies. Figure 17 plots default rates in our data by the origination risk score categories we use, ranging from the least risky (760+) to the riskiest deep subprime borrowers (< 560). While default rates are generally near zero for the least risky borrowers, they often increase more than 5 percentage points for the riskiest borrowers. Figure 18 plots default rates for loans originated by the lender categories that we observe in our data. Auto finance loans typically have default rates above 1.5 percent, and their default rates are usually at least a full percentage point higher than those of banks or credit unions.

To highlight where the default forecast errors are largest using the model coefficients as of December 2019, we examine forecast errors in both absolute differences and as a percent of the actual default rate. Figure 19 plots the cumulative default transition forecast differences relative to actual outcomes by original risk score groupings. In absolute terms, the forecast errors are largest for the riskiest borrowers, but in percent terms the forecast errors are largest for the least risky borrowers. The absolute difference is reflective of the higher default rates for low-quality borrowers. Small errors, however, can be large in percentage terms for the highest quality borrowers because they have very low default rates.

Figure 20 plots the same analysis of cumulative forecast differences for origination lender types. Similar to the risk score results, the absolute errors are largest for auto finance lenders, which originate loans to a larger share of risky borrowers. The percentage difference is relatively higher for banks and credit unions, which hold less risk in their auto loan portfolios and face lower default rates.

Stress testing is concerned with forecasted losses in levels so inaccuracies in level forecasts can be economically meaningful for lenders or regulators in terms of capital levels. Alternatively, percentage forecast errors relative to realized default rates highlight model calibration issues in the statistical sense, that is, whether predicted probabilities match actual default frequencies across the risk distribution. For example, because default rates among the safest borrowers are low, small deviations can lead to large percentage errors, mischaracterizing the macroeconomic sensitivity of safer borrowers and distorting risk differentiation across the credit risk distribution. The percentage forecast errors we find are notable because the majority of auto lending is made to higher quality borrowers, so that small absolute errors can aggregate to large errors for an overall portfolio of loans. This can have large economic implications for loss planning for lenders in addition to the large absolute errors for the riskiest segments of the loan population.

8.2 Oaxaca-Blinder Decomposition

While we find evidence that the statistical relationship between unemployment and default changed in 2020 and that this change explains the large forecast error, it is possible that the forecast error is also due to changes in borrower and loan characteristics. To evaluate this possibility, we follow Fairlie (2005) and use a nonlinear Oaxaca-Blinder decomposition on our data and rolling coefficient estimates. This procedure involves computing average predicted probabilities from the data and model coefficients in two different periods, a baseline period and an alternative period. The total difference in the average predicted probabilities over the two periods is then decomposed into components explained by the changes in loan data, macroeconomic data, and the estimated coefficients.

We set December 2019 as the baseline period and December 2021 as the alternative period. The predictions use various combinations of the borrower and loan characteristics $\boldsymbol{x}$, macroeconomic variables $\boldsymbol{z}$ and the estimated borrower and loan coefficients $\hat{\boldsymbol{\beta}}$ and the estimated macro shock coefficients $\hat{\boldsymbol{\gamma}}$ from each of the two time periods. Taking our model's average predicted probabilities for transitions from state j to state k as $\overline{p}_{jk}$, define the total change in average predicted probabilities between the baseline 2019 period b and alternative

2021 period a

$$\Delta^{Total} = \overline{p}_{jk}(\boldsymbol{x_a}, \boldsymbol{z_a}; \hat{\boldsymbol{\beta}}_{\boldsymbol{a}}, \hat{\boldsymbol{\gamma}}_{\boldsymbol{a}}) - \overline{p}_{jk}(\boldsymbol{x_b}, \boldsymbol{z_b}; \hat{\boldsymbol{\beta}}_{\boldsymbol{b}}, \hat{\boldsymbol{\gamma}}_{\boldsymbol{b}}) \tag{6}$$

$$\Delta^{Borrower} = \overline{p}_{jk}(\boldsymbol{x_a}, \boldsymbol{z_b}; \hat{\boldsymbol{\beta}}_{\boldsymbol{a}}, \hat{\boldsymbol{\gamma}}_{\boldsymbol{a}}) - \overline{p}_{jk}(\boldsymbol{x_b}, \boldsymbol{z_b}; \hat{\boldsymbol{\beta}}_{\boldsymbol{a}}, \hat{\boldsymbol{\gamma}}_{\boldsymbol{a}}) \tag{7}$$

$$\Delta^{Macro} = \overline{p}_{jk}(\boldsymbol{x_a}, \boldsymbol{z_a}; \hat{\boldsymbol{\beta}}_{\boldsymbol{a}}, \hat{\boldsymbol{\gamma}}_{\boldsymbol{a}}) - \overline{p}_{jk}(\boldsymbol{x_a}, \boldsymbol{z_b}; \hat{\boldsymbol{\beta}}_{\boldsymbol{a}}, \hat{\boldsymbol{\gamma}}_{\boldsymbol{a}}) \tag{8}$$

$$\Delta^{Coefficient} = \overline{p}_{jk}(\boldsymbol{x_b}, \boldsymbol{z_b}; \hat{\boldsymbol{\beta}}_{\boldsymbol{a}}, \hat{\boldsymbol{\gamma}}_{\boldsymbol{a}}) - \overline{p}_{jk}(\boldsymbol{x_b}, \boldsymbol{z_b}; \hat{\boldsymbol{\beta}}_{\boldsymbol{b}}, \hat{\boldsymbol{\gamma}}_{\boldsymbol{b}}) \tag{9}$$

This specification examines what explains the difference between predictions for 2021 made in 2019 and the in-sample prediction for 2021 made using the rolling regression estimated through 2021. The coefficient effect measures how much the model change matters when applied to 2019 conditions, while the borrower and macro effects measure how much data changes matter when evaluated with the new model. In this decomposition, coefficients are updated holding the data fixed at the baseline. Then, holding coefficients fixed at the alternative levels, borrower composition is updated holding macro variables fixed at baseline levels and macro variables are updated holding borrower characteristics fixed at alternative levels. This allows the total predicted change to be decomposed as the sum of the components,

$$\Delta^{Total} = \Delta^{Borrower} + \Delta^{Macro} + \Delta^{Coefficient} \tag{10}$$

To examine the relative contributions to the total effects, Figure 21 plots the percent of the total effect explained by each component for the two absorbing states. The results show that macroeconomic changes dominate (60-70 percent), reflecting the massive unemployment spike from 3 to 13 percent. Even with the weakened unemployment sensitivity captured by 2021 coefficients, this shock is so large it overwhelms other factors. The coefficient effect (20-30 percent) captures the model change, but is smaller because it is evaluated on 2019 data when unemployment is low. The borrower composition effect is much smaller than the macroeconomic effect (18 percent), ruling out composition changes as the primary driver of the forecast errors.[13]

The decomposition results are consistent with a large macroeconomic shock that generates a large forecast error rather than the error being meaningfully impacted by changes in borrower composition. Taken together, our results show that the resulting forecast errors are generated by a large change in the macroeconomic data rather than a large change in the underlying population of borrowers.

8.3 Comparison to Binomial Discrete-Time Hazard Model

A simpler approach to estimating probabilities of default for discrete-time survival analysis or hazard modeling is to use a single binary outcome regression. Indeed, papers in the academic literature as well as regulatory stress testing models rely on this approach rather than constructing an entire transition matrix. The binomial discrete-time hazard model estimates the probability of default without taking into account probabilistic movement throughout the state space and the competing risk of non-default exit. The binomial approach may be appropriate when loan dynamics are relatively simple—for instance, when exit rates are low, cure rates are minimal, and state transitions are rare. However, when these simplifying assumptions are violated, the binomial model can produce biased forecasts.

Our historical transition rates across states motivate why we might care about the transition approach. On average, there is a 40 percent probability that loans in our sample that start as delinquent will "cure" to current, and more than a 12 percent probability that loans starting as current will exit without default. Failing to take these probabilities into account when estimating a cumulative probability can impact the accuracy of the estimates, especially in multi-period-ahead forecasts.

For comparison, we estimate a binomial discrete-time hazard model. The binomial model

[13] Given that the effects are nonlinear, Fairlie (2005) notes that the ordering of the effects matters. An alternative ordering of the effects produces nearly identical results.

takes the form

$$\ln\left(\frac{p_{i,df}(t)}{1-p_{i,df}(t)}\right) = \boldsymbol{x_{it}\beta} + \boldsymbol{z_t\gamma} \tag{11}$$

where $p_{i,df}(t)$ represents the 1-period-ahead probability of default for loan i at time t. Solving for $p_{i,df}$, the h-period-ahead cumulative probability of default P_{it}^h, accounting for the probability of surviving each previous period, can be recursively constructed as

$$\begin{aligned}
P_{it}^1 &= p_{i,df}(t) \\
P_{it}^2 &= P_{it}^1 + (1 - P_{it}^1) \cdot p_{i,df}(t+1) \\
&\vdots \\
P_{it}^h &= P_{it}^{h-1} + (1 - P_{it}^{h-1}) \cdot p_{i,df}(t+h-1)
\end{aligned}$$

Note that this specification accounts for survival (non-default) but not for exits or state transitions. A loan that exits in period $t = 2$ continues to enter the survival probability calculation in subsequent periods, and a delinquent loan that cures continues to face delinquent-state default probabilities. This contrasts with the transition matrix approach, where the state vector $\boldsymbol{P}(\boldsymbol{x_{it}}, \boldsymbol{z_t})$ dynamically adjusts to reflect exits, cures, and risk transitions.

We estimate (11) using our data from 2000-2019 and use it to produce two-year cumulative forecasts over the period 2020-2021. As previously shown in Figure 4, actual transitions to default declined during that period, but transitions from current to exit and from delinquent to current both increased. Figure 22 plots the path of the cumulative default predictions implied by the binomial logit model and the Markov transition and the actual cumulative default outcomes over the 2020-2021 period. The initial forecast in June 2020 is roughly equivalent between the two models. However, by the end of the forecast period, the binomial model produces a cumulative forecast error that is 34 percent higher than the actual cumulative outcome while the Markov transition matrix is 24 percent higher.

While the Markov model produces lower cumulative forecast errors during the 2020-2021 period, it is instructive to better understand if the Markov transition model may be more accurate in capturing the dynamics of borrower transitions across other periods in our sample. We use both models to produce a two-year cumulative forecast at each semiannual period in our sample, allowing the macroeconomic variables to evolve as they actually did and treating the loan characteristics as previously described. Table 5 shows the total cumulative forecast error comparison between the Markov transition matrix and the binomial model across all time periods. The Markov transition model has lower cumulative forecast errors across standard measures that include the mean absolute error (MAE), the mean absolute percentage error (MAPE), and root mean squared error (RMSE). Figure 23 plots the rolling two-year cumulative default forecast errors for the Markov transition matrix and binomial logit model at each date in our sample. The Markov transition model produces lower cumulative forecast errors in almost all periods, with the exception of several years from 2015-2017 when the exit transition forecasts deteriorated and in 2025.

These results imply that the Markov transition model's estimated overprediction of cumulative default rates using estimated model parameters as of 2019 are conservative relative to simpler modeling approaches that may be used by lenders and regulators. It also implies that the forecast errors in the binomial model will be driven not only by the unstable sensitivity of default rates to unemployment, but also by the fact that the binomial model does not account for all of the dynamic transitions across the state space.

9 Conclusion

We estimate a Markov transition model of auto loan performance over the period from 2000-2025 using a panel of semiannual loan-level data. The pandemic period was characterized by large changes in macroeconomic variables typically used to model loan performance under stress, but loan performance did not deteriorate in 2020-2021 in part due to large government transfers that offset the shocks. We examined the implications for loss forecasting using our model to produce forecasts in a rolling pseudo out-of-sample framework. We also document a lender channel of loan performance risk, where the riskiest loans are originated by the auto finance industry.

Using the Markov transition framework, we produce multi-period-ahead forecasts that take into account the probabilistic movements between states of loan performance. We find large single and multi-period-ahead default forecast errors during the pandemic period and instability in the estimated parameters on unemployment. The forecast errors are attributable the instability of the unemployment coefficient in the model that occurs after large aggregate shocks, first during the 2007-2009 recession and again during the 2020 Covid recession. Moreover we showed that the instability leads to persistently lower levels of sensitivity to unemployment shocks for loans that transition to default rather than eventually returning to prior levels. Counterfactual forecasts and Oaxaca-Blinder decompositions showed that the forecast errors and coefficient instability were primarily driven by the model learning the revised relationships between unemployment and loan performance outcomes. We show that forbearance does not appear to explain the unstable relationship during Covid, and robustness tests that attempt to control for the 2020-2021 period using dummy variables do not improve the forecast errors or parameter instability.

The results suggest that the relationship between unemployment and auto loan default is declining to a point where it is nearly meaningless for loan default predictions. If our model were used to evaluate policies around auto loan defaults, there are obvious connections to the Lucas critique (Lucas (1976)) given the reduced-form relationships used to model the impact of macroeconomic outcomes on aggregate loan performance. Given how common unemployment is in reduced-form models of consumer loan loss forecasting, these findings are concerning for lenders and regulators that rely on these estimated relationships to correctly forecast future losses. Future work could examine how structural approaches to loan loss modeling may improve on forecasts of loan performance in periods where government interventions or changes in lending practices due to a recession change statistical relationships.

References

Agarwal, Sumit (2008). “Determinants of Automobile Loan Default and Prepayment.” *Economic Perspectives* 32 (3).

Agarwal, Sumit and Brent W. Ambrose, eds. (2007). *Household Credit Usage*. New York: Palgrave Macmillan US. DOI: 10.1057/9780230608917.

An, Xudong, Larry Cordell, and Sharon Tang (2020). “Extended Loan Terms and Auto Loan Default Risk.” *Federal Reserve Bank of Philadelphia Working Paper* 20-18. DOI: 10.21799/frbp.wp.2020.18.

Bakshi, Tanya and Jonathan Rose (2021). “What Happened to Subprime Auto Loans During the COVID-19 Pandemic?” *Chicago Fed Insights*.

Canals-Cerda, José J. and Brian Jonghwan Lee (2021). “COVID-19 and Auto Loan Origination Trends.” *SSRN Electronic Journal*. DOI: 10.2139/ssrn.3897908.

Chen, Qingqing, Dennis Glennon, and Amos Golan (2020). “Estimating Conditional Mortgage Delinquency Transition Matrices.” *OCC Working Paper Series* 2020-05.

Covas, Francisco B., Ben Rump, and Egon Zakrajšek (July 2014). “Stress-Testing US Bank Holding Companies: A Dynamic Panel Quantile Regression Approach.” *International Journal of Forecasting* 30 (3), pp. 691–713. DOI: 10.1016/j.ijforecast.2013.11.003.

Deng, Yongheng, John M. Quigley, and Robert Order (2000). “Mortgage Terminations, Heterogeneity and the Exercise of Mortgage Options.” *Econometrica* 68 (2), pp. 275–307. DOI: 10.1111/1468-0262.00110.

Dettling, Lisa and Lauren Lambie-Hanson (2021). “Why is the Default Rate So Low? How Economic Conditions and Public Policies Have Shaped Mortgage and Auto Delinquencies

During the COVID-19 Pandemic." *Board of Governors of the Federal Reserve System FEDS Notes,* March 4, 2021. DOI: 10.17016/2380-7172.2854.

Fairlie, Robert W. (2005). "An Extension of the Blinder-Oaxaca Decomposition Technique to Logit and Probit Models." *Journal of Economic and Social Measurement* 30 (4), pp. 305–316. DOI: 10.3233/JEM-2005-0259.

Gross, Marco and Javier Población (Feb. 2019). "Implications of Model Uncertainty for Bank Stress Testing." *Journal of Financial Services Research* 55 (1), pp. 31–58. DOI: 10.1007/s10693-017-0275-4.

Guerrieri, Luca and Michelle Welch (2012). "Can Macro Variables Used in Stress Testing Forecast the Performance of Banks?" *Finance and Economics Discussion Series* 2012-49. DOI: 10.17016/feds.2012.49.

Guo, Zhengfeng, Yan Zhang, and Xinlei Zhao (2022). "Risks of Long-Term Auto Loans." *Journal of Credit Risk* 18 (4). DOI: 10.21314/JCR.2022.009.

Heitfield, Erik and Tarun Sabarwal (2004). "What Drives Default and Prepayment on Subprime Auto Loans?" *The Journal of Real Estate Finance and Economics* 29 (4), pp. 457–477. DOI: 10.1023/B:REAL.0000044023.02636.e6.

Hirtle, Beverly, Anna Kovner, James Vickery, and Meru Bhanot (Aug. 2016). "Assessing Financial Stability: The Capital and Loss Assessment under Stress Scenarios (CLASS) Model." *Journal of Banking & Finance* 69, S35–S55. DOI: 10.1016/j.jbankfin.2015.09.021.

Katcher, Bradley, Geng Li, Alvaro Mezza, and Steve Ramos (2024). "One Month Longer, One Month Later? Prepayments in the Auto Loan Market." *Finance and Economics Discussion Series* 2024-056. DOI: 10.17016/feds.2024.056.

Lancaster, Tony (1990). *The Econometric Analysis of Transition Data.* Cambridge University Press.

Lucas, Robert E. (1976). "Econometric Policy Evaluation: A Critique." *Carnegie-Rochester Conference Series on Public Policy* 1, pp. 19–46. DOI: 10.1016/S0167-2231(76)80003-6.

Smith, Douglas L. and Edward C. Lawrence (1995). "Forecasting Losses on a Liquidating Long-Term Loan Portfolio." *Journal of Banking & Finance* 19 (6), pp. 959–985. DOI: 10.1016/0378-4266(94)00065-B.

Smith, L. Douglas and Baiqiang Jin (2007). "Modeling Exposure to Losses on Automobile Leases." *Review of Quantitative Finance and Accounting* 29 (3), pp. 241–266. DOI: 10.1007/s11156-007-0032-0.

Smith, L. Douglas, Susan M. Sanchez, and Edward C. Lawrence (1996). "A Comprehensive Model for Managing Credit Risk on Home Mortgage Portfolios." *Decision Sciences* 27 (2), pp. 291–317. DOI: 10.1111/j.1540-5915.1996.tb01719.x.

Wu, Deming, Ming Fang, and Qing Wang (2018). "An Empirical Study of Bank Stress Testing for Auto Loans." *Journal of Financial Stability* 39, pp. 79–89. DOI: 10.1016/j.jfs.2018.09.005.

Table 1: Analysis variables used in multinomial logit.

Variable	Description	Source
Loan Age	Loan age in months.	FRBNY Consumer Credit Panel / Equifax Data
Loan Age (squared)	Loan age in months (squared).	FRBNY Consumer Credit Panel / Equifax Data
Risk Score	Equifax risk score group at origination. One of <560, 560-619, 620-659, 660-719, 720-760, 760+	FRBNY Consumer Credit Panel / Equifax Data
Lender	Lender type. One of auto finance, bank, or credit union.	FRBNY Consumer Credit Panel / Equifax Data
Borrower Age	Borrower age (years).	FRBNY Consumer Credit Panel / Equifax Data
Joint Account	Equal to 1 if tradeline is a joint account, 0 otherwise.	FRBNY Consumer Credit Panel / Equifax Data
First Mortgage Flag (t-1)	Equal to 1 if borrower is a first mortgage holder as of the prior period, 0 otherwise.	FRBNY Consumer Credit Panel / Equifax Data
Ln(Balance) (t-1)	Natural log of outstanding balance in the prior period.	FRBNY Consumer Credit Panel / Equifax Data
Annual Unemployment Change	Annual change in US state unemployment.	BLS
Annual HPI Change	Annual change in US state House Price Index	FHFA
Seasonal Flag	Equal to 1 if period is in December-May, 0 otherwise.	FRBNY Consumer Credit Panel / Equifax Data

Table 2: Summary statistics of analysis variables, biannual observations from 2000–2025.

Variable	N	Mean	SD	Min	P50	Max
Current	1668736.00	0.82	0.38	0.00	1.00	1.00
30-89 DPD	1668736.00	0.03	0.16	0.00	0.00	1.00
90+ DPD	1668736.00	0.01	0.11	0.00	0.00	1.00
Non-Default Exit	1668736.00	0.14	0.34	0.00	0.00	1.00
Loan Age (months)	1668736.00	24.51	17.49	0.00	21.00	108.00
Equifax Risk Score	1668736.00	690.69	104.37	280.00	710.00	841.00
Borrower Age	1668736.00	45.97	14.34	17.00	45.00	99.00
Joint Account	1668736.00	0.43	0.50	0.00	0.00	1.00
Remaining Balance	1668736.00	11005.85	7234.20	0.52	9704.76	129598.15
First Mortgage Flag	1668736.00	0.31	0.46	0.00	0.00	1.00
Annual Change in State Unemployment (p.p.)	1668736.00	-0.04	2.26	-17.47	-0.27	20.63
Annual State House Price Index Change (p.p.)	1668736.00	4.65	6.65	-30.16	4.68	30.54

Source: FRBNY Consumer Credit Panel / Equifax Data, FHFA, BLS, Authors' calculations.

Table 3: Current State Multinomial Logit Results

	y = Transitions from Current State		
	Current-Delinquent	Current-Exit	Current-Default
Intercept	0.000***	306.610***	0.032***
	[0.000, 0.000]	[306.531, 306.689]	[0.032, 0.032]
Lender: Auto Finance	1.705***	0.967***	1.551***
	[1.694, 1.716]	[0.954, 0.980]	[1.550, 1.553]
Lender: Banks	1.037***	1.016**	0.886***
	[1.033, 1.041]	[1.001, 1.032]	[0.885, 0.886]
Loan Age (months)	1.041***	1.004***	1.037***
	[1.038, 1.044]	[1.003, 1.005]	[1.033, 1.041]
Loan Age ^{2}	1.000***	1.000***	1.000***
	[1.000, 1.000]	[1.000, 1.000]	[1.000, 1.000]
Equifax Risk Score: 720-759	3.003***	0.951***	1.553***
	[3.002, 3.004]	[0.938, 0.964]	[1.553, 1.553]
Equifax Risk Score: 660-719	7.572***	0.940***	2.735***
	[7.559, 7.585]	[0.929, 0.951]	[2.735, 2.736]
Equifax Risk Score: 620-659	17.505***	0.953***	4.905***
	[17.464, 17.546]	[0.940, 0.966]	[4.904, 4.907]
Equifax Risk Score: <560	69.243***	0.827***	22.227***
	[68.077, 70.429]	[0.815, 0.839]	[22.188, 22.265]
Borrower Age	0.995***	0.997***	0.993***
	[0.994, 0.996]	[0.996, 0.997]	[0.992, 0.995]
First Mortgage Flag (t-1)	0.845***	1.070***	0.687***
	[0.843, 0.846]	[1.058, 1.083]	[0.687, 0.687]
Ln(Balance) (t-1)	1.277***	0.438***	0.721***
	[1.269, 1.284]	[0.437, 0.439]	[0.715, 0.727]
Joint Account	0.917***	0.961***	0.868***
	[0.916, 0.917]	[0.950, 0.972]	[0.868, 0.868]
Annual Unemployment Change	1.008	0.998	1.089***
	[0.995, 1.021]	[0.993, 1.003]	[1.070, 1.109]
Annual HPI Change	0.990***	1.012***	1.003**
	[0.988, 0.992]	[1.011, 1.013]	[1.000, 1.007]
Seasonal Effect (Jan-June = 1)	0.821***	1.011**	0.788***
	[0.821, 0.822]	[1.001, 1.022]	[0.788, 0.789]
Num.Obs.	1181077		
R2	0.110		
R2 Adj.	0.110		
Pseudo-R2	0.11		
AUC	0.71		

Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations. Model estimated on sample data from 2000-2019. Coefficients represent relative risks. *, **, *** represent statistical significance at 0.1, 0.05, and 0.01 levels, respectively. 95 percent confidence intervals shown in brackets. Standard errors are clustered at the tradeline level.

Table 4: Delinquent State Multinomial Logit Results

	y = Transitions from Delinquent State		
	Delinquent-Delinquent	Delinquent-Exit	Delinquent-Default
Intercept	0.263***	379.628***	133.403***
	[0.262, 0.263]	[378.849, 380.410]	[133.181, 133.627]
Lender: Auto Finance	1.086***	0.818***	1.096***
	[1.045, 1.129]	[0.789, 0.847]	[1.050, 1.144]
Lender: Banks	0.866***	1.113***	0.929***
	[0.842, 0.891]	[1.083, 1.144]	[0.901, 0.957]
Loan Age (months)	0.996	0.987***	0.942***
	[0.989, 1.003]	[0.977, 0.996]	[0.935, 0.949]
Loan Age ^{2}	1.000***	1.000***	1.001***
	[1.000, 1.000]	[1.000, 1.001]	[1.001, 1.001]
Equifax Risk Score: 720-759	1.565***	0.757***	1.170***
	[1.560, 1.569]	[0.755, 0.758]	[1.168, 1.172]
Equifax Risk Score: 660-719	1.559***	0.889***	1.298***
	[1.467, 1.658]	[0.873, 0.905]	[1.269, 1.329]
Equifax Risk Score: 620-659	2.080***	0.912***	1.634***
	[1.977, 2.190]	[0.882, 0.943]	[1.561, 1.711]
Equifax Risk Score: <560	2.843***	0.862***	2.542***
	[2.722, 2.969]	[0.816, 0.912]	[2.428, 2.662]
Borrower Age	1.002**	1.000	0.998
	[1.000, 1.005]	[0.997, 1.004]	[0.996, 1.001]
First Mortgage Flag (t-1)	0.875***	0.991	0.541***
	[0.822, 0.930]	[0.975, 1.008]	[0.528, 0.555]
Ln(Balance) (t-1)	1.013	0.453***	0.601***
	[0.997, 1.030]	[0.442, 0.464]	[0.590, 0.611]
Joint Account	0.958*	0.925***	0.925***
	[0.914, 1.004]	[0.896, 0.956]	[0.878, 0.975]
Annual Unemployment Change	1.008	1.057***	0.994
	[0.982, 1.034]	[1.017, 1.099]	[0.964, 1.023]
Annual HPI Change	1.001	1.021***	0.989***
	[0.996, 1.006]	[1.013, 1.029]	[0.984, 0.995]
Seasonal Effect (Jan-June = 1)	0.860***	0.882***	0.753***
	[0.821, 0.900]	[0.854, 0.910]	[0.715, 0.793]
Num.Obs.	29835		
R2	0.048		
R2 Adj.	0.048		
Pseudo-R2	0.05		
AUC	0.65		

Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations. Model estimated on sample data from 2000-2019. Coefficients represent relative risks. *, **, *** represent statistical significance at 0.1, 0.05, and 0.01 levels, respectively. 95 percent confidence intervals shown in brackets. Standard errors are clustered at the tradeline level.

Table 5: Comparison of pseudo out-of-sample 2-year cumulative forecast errors across 2007-2025.

Model	MAE	MAPE	RMSE
Binomial Logit	0.79	18	1
Markov Transition	0.44	9.9	0.5

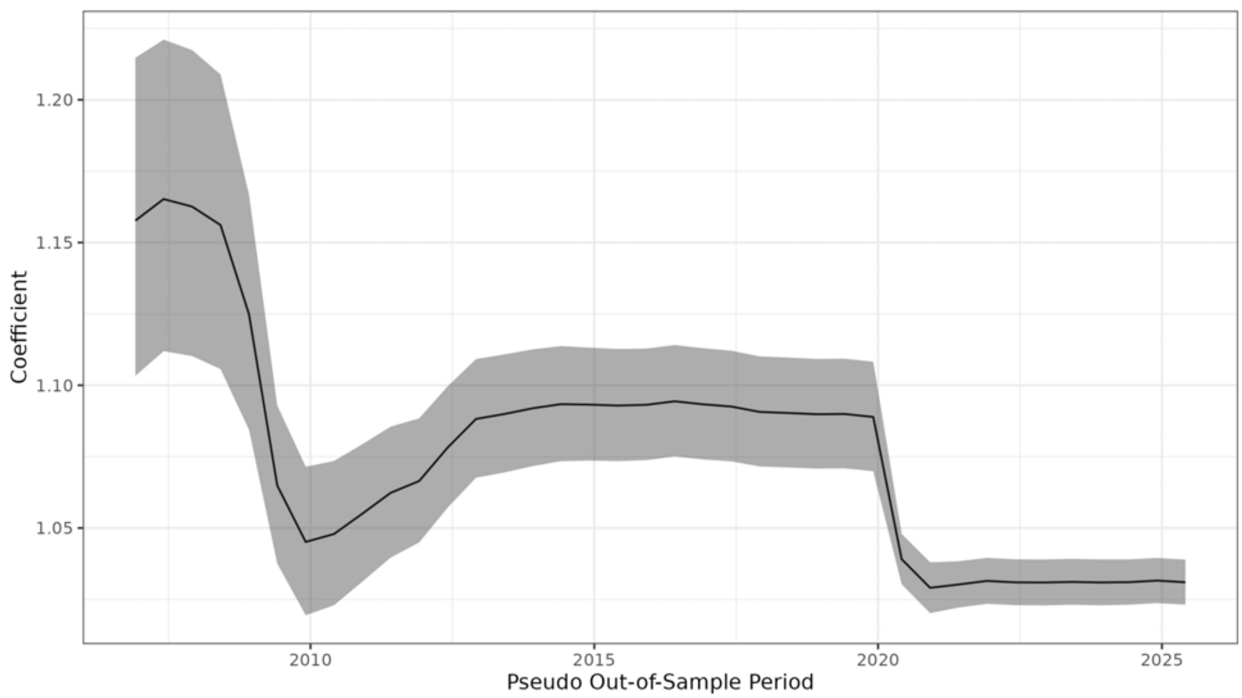

Figure 1: Multinomial logit estimates of US state unemployment coefficients for auto loan transitions from current to default in pseudo out-of-sample estimations, 2007-2025. Estimation variables include loan age, lender type, loan characteristics, and year-over-year changes in US state unemployment and HPI. Shaded areas represent 95 percent confidence intervals, where standard errors are clustered at the loan level. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

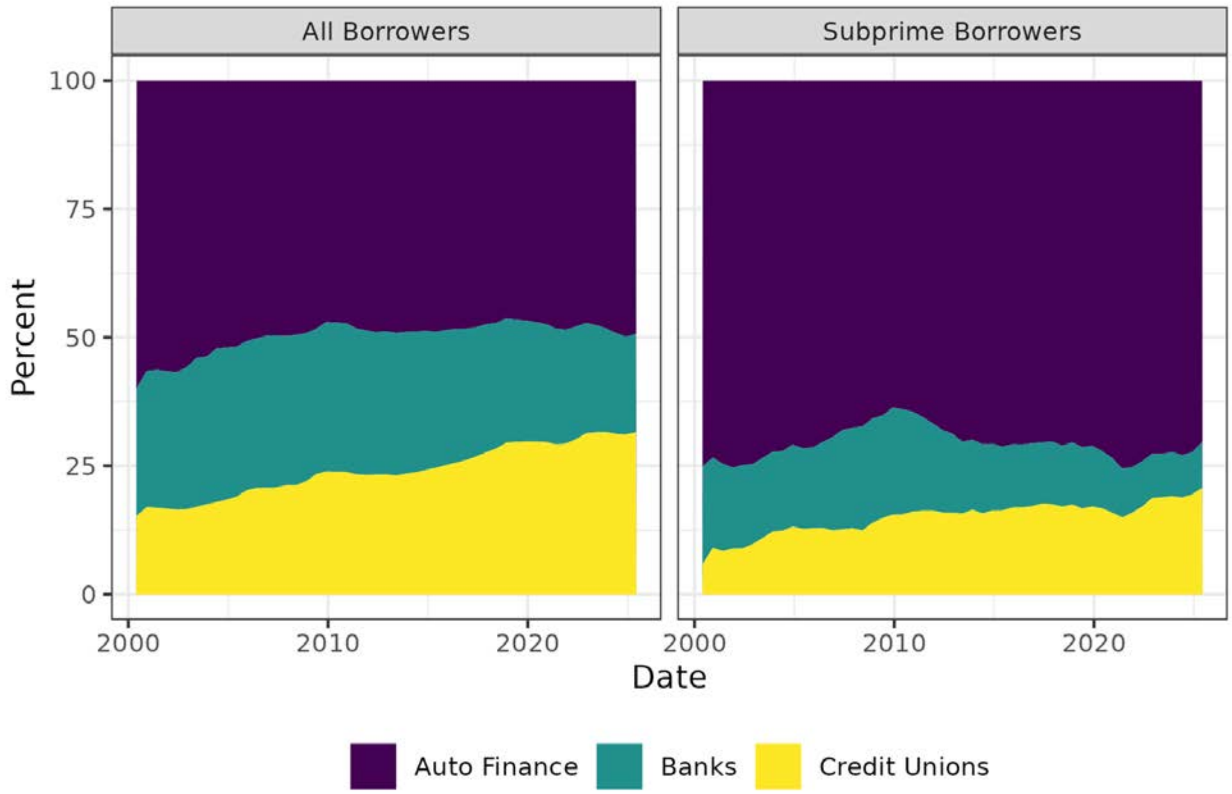

Figure 2: Lender types as a percent of total outstanding loans. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

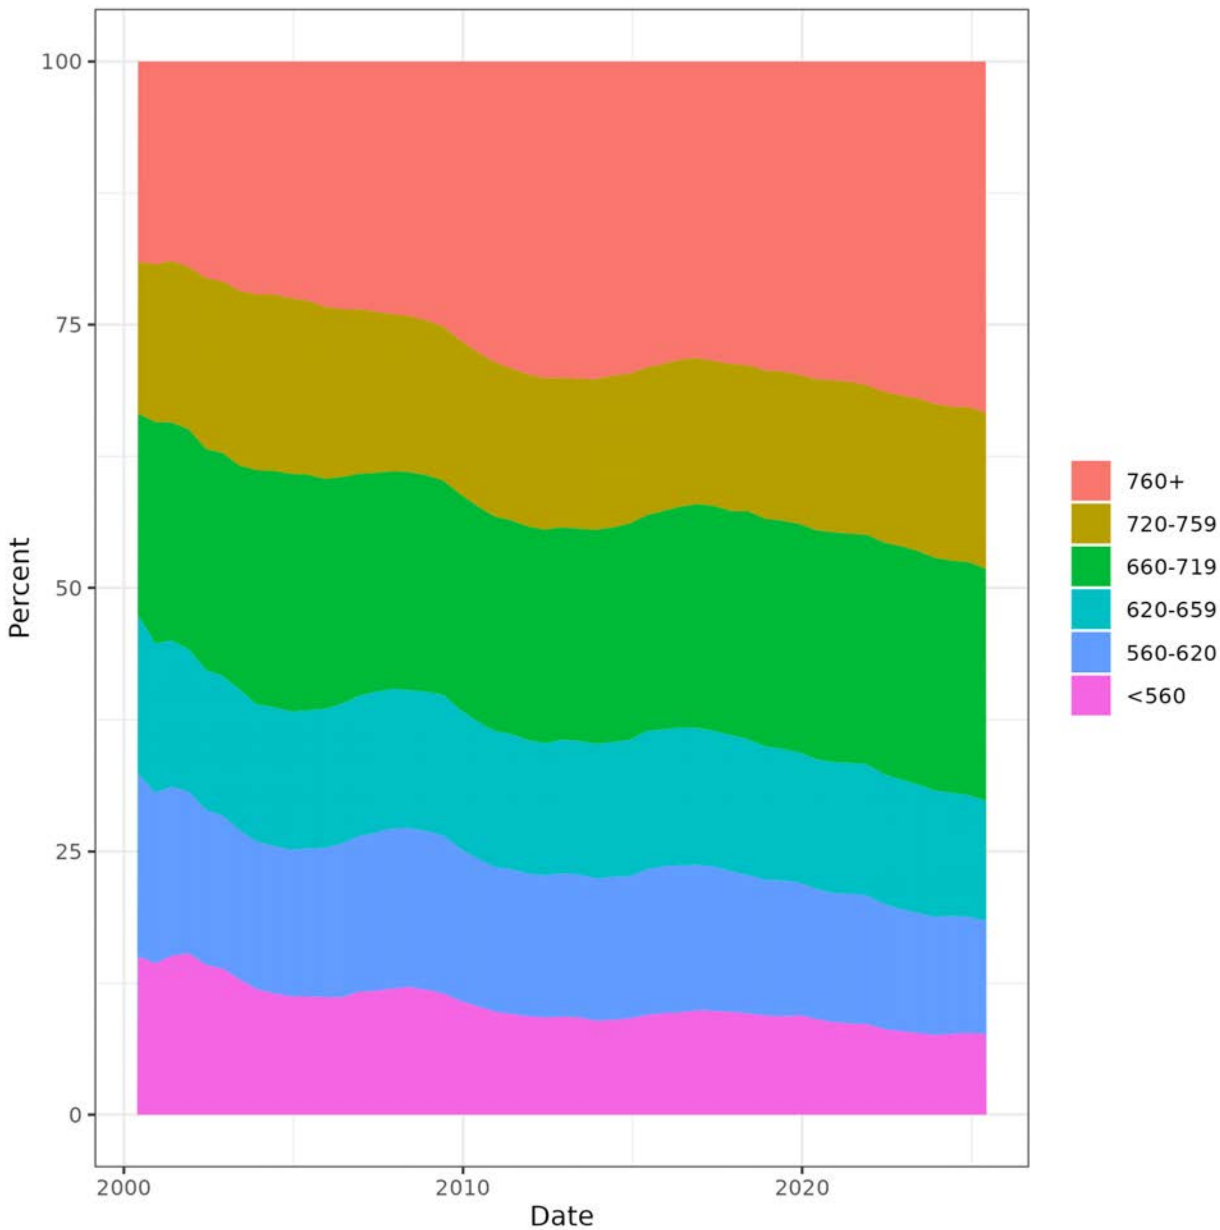

Figure 3: Origination Equifax Risk Scores as a percent of total outstanding loans. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

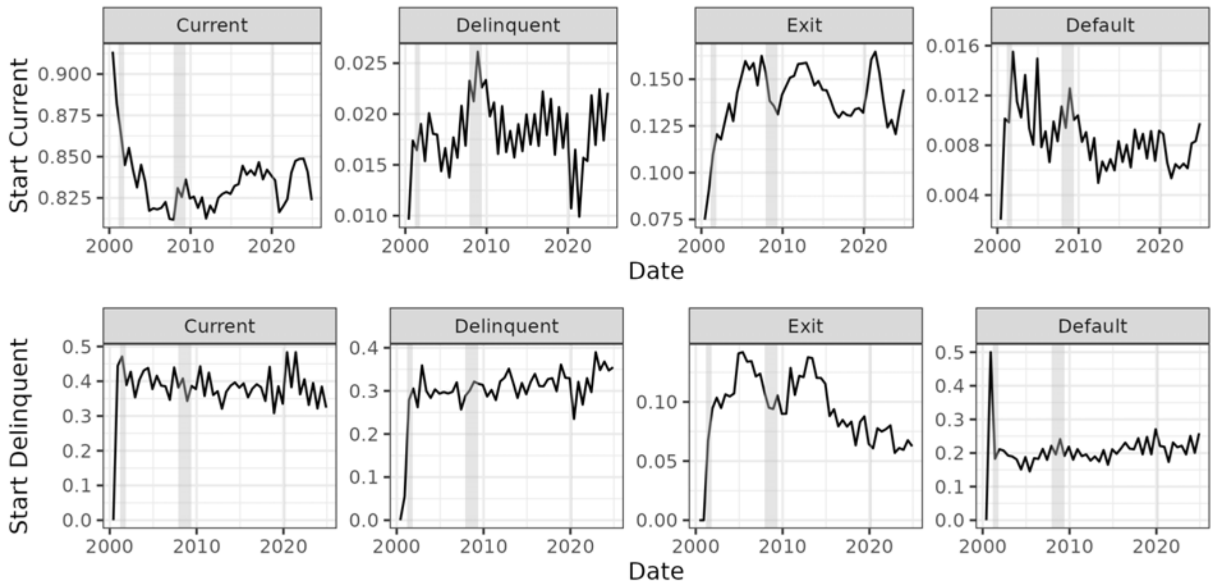

Figure 4: Actual transition rates by starting state, December 2000-June 2025. Note: transition rates for first few periods are distorted due to censoring at the beginning of the dataset. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

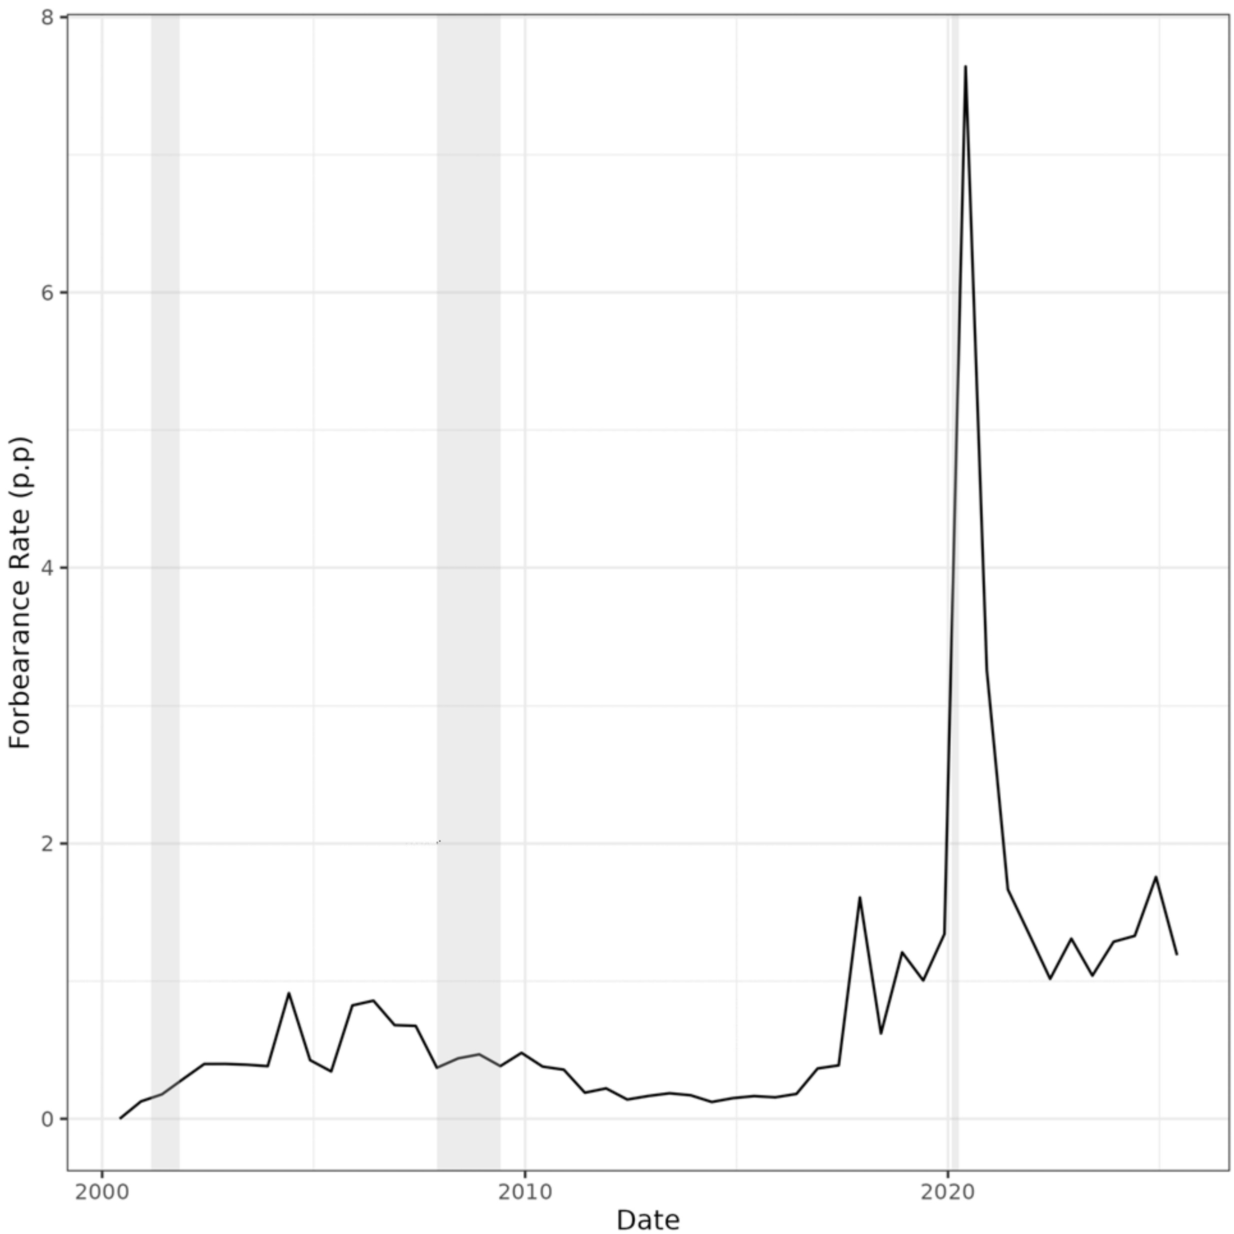

Figure 5: Auto loan forbearance usage, semiannual data from 2000-June 2025. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

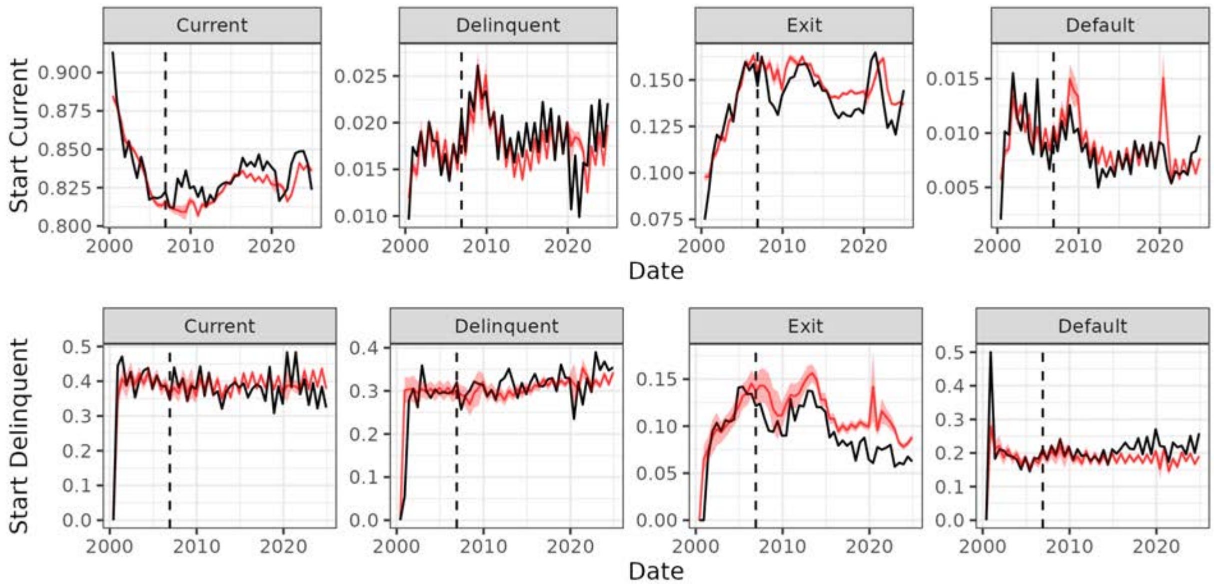

Figure 6: Time path of current and delinquent-state transition probabilities from pseudo out-of-sample estimation. Estimation variables include loan age, lender type, loan characteristics, and year-over-year changes in US state unemployment and HPI. The red line represents pseudo out-of-sample forecasts and the black line represents actual outcomes. Red shaded areas represent 95 percent confidence intervals, where standard errors are clustered at the loan level. Dashed line indicates December 2006, when the rolling regressions start. Pre-December 2006 estimates are in sample. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

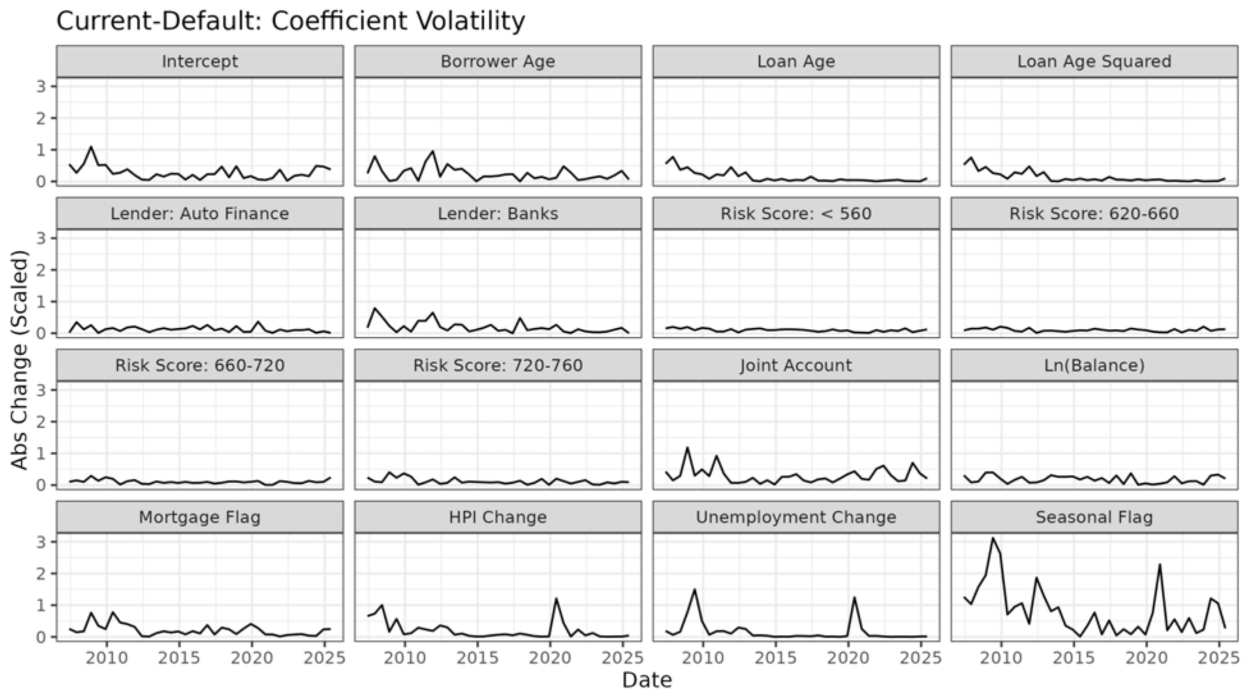

Figure 7: Rolling standard deviations of semiannual change in estimated relative risk coefficients for transitions from current to default. Note: Risk Score refers to the Equifax Risk Score. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

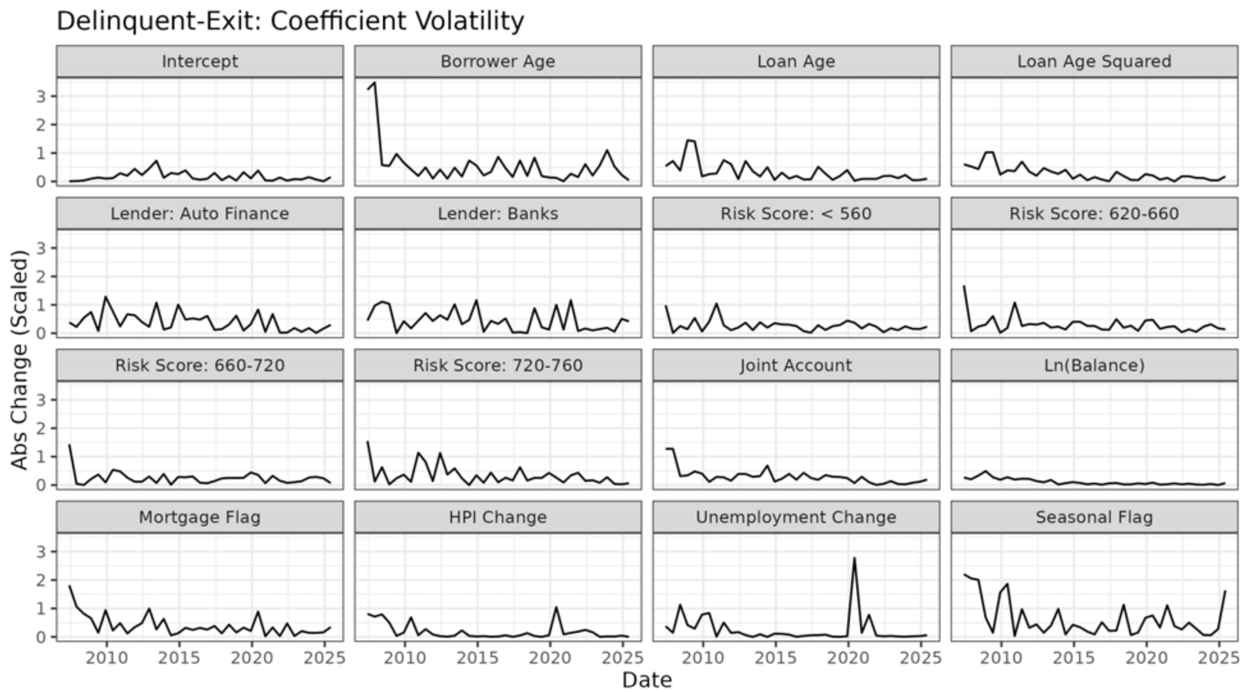

Figure 8: Rolling standard deviations of semiannual change in estimated relative risk coefficients for transitions from delinquent to exit. Note: Risk Score refers to the Equifax Risk Score. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

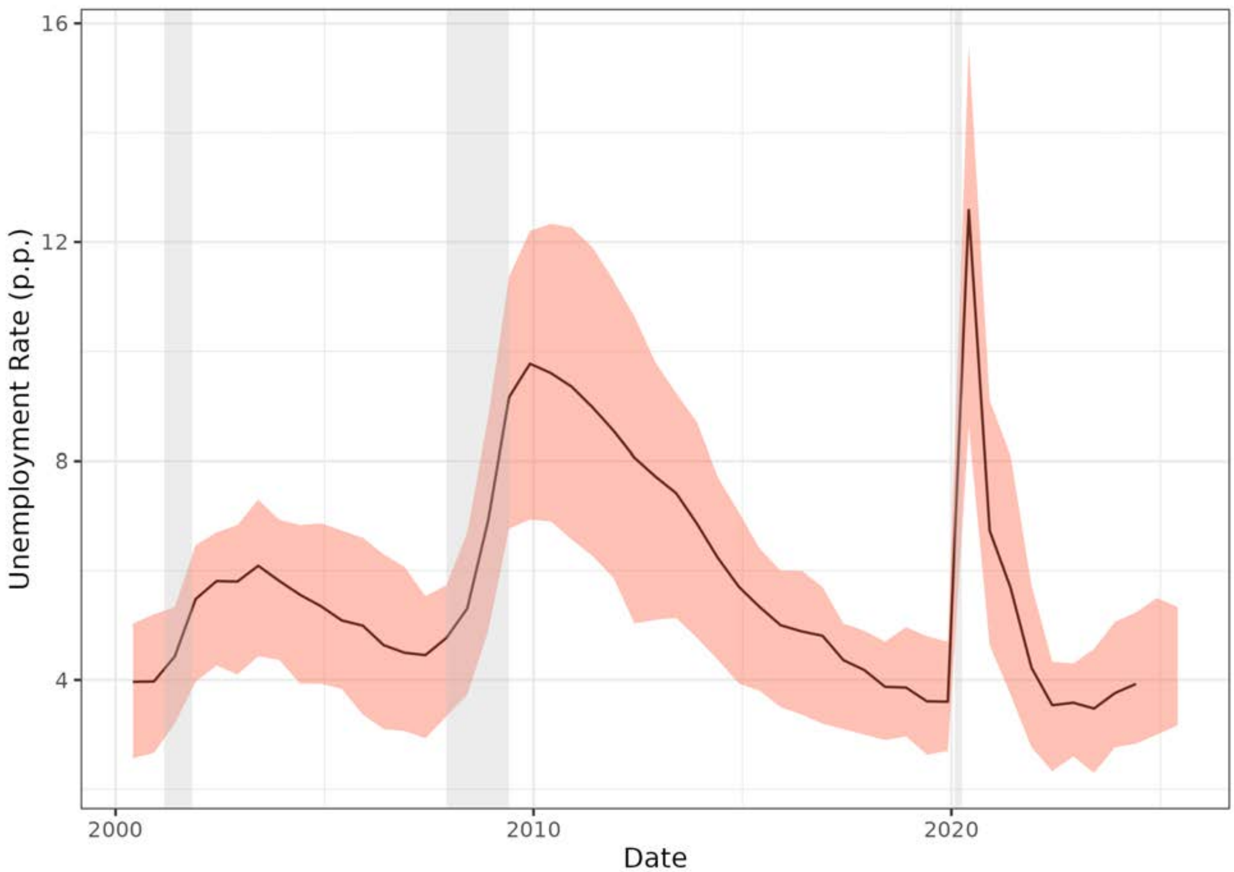

Figure 9: US state unemployment; semiannual data from 2000-June 2025. The solid line represents the average; the shaded area represents the 5th-95th percentiles. Source: BLS

Figure 10: US state house price index; semiannual data from 2000-June 2025. The solid line represents the average; the shaded area represents the 5th-95th percentiles. Source: FHFA

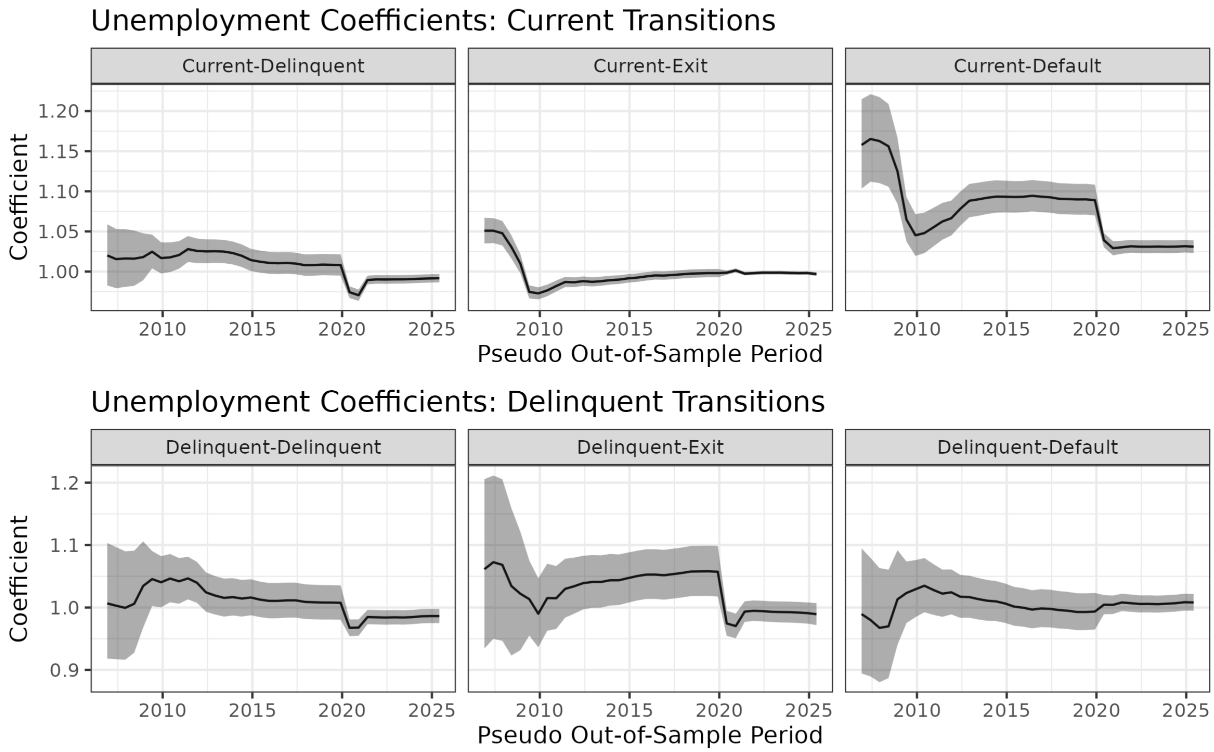

Figure 11: Multinomial logit estimates of US state unemployment coefficients during pseudo out-of-sample estimations, 2007-2025. Estimation variables include loan age, lender type, loan characteristics, and year-over-year changes in US state unemployment and HPI. Shaded areas represent 95 percent confidence intervals, where standard errors are clustered at the loan level. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

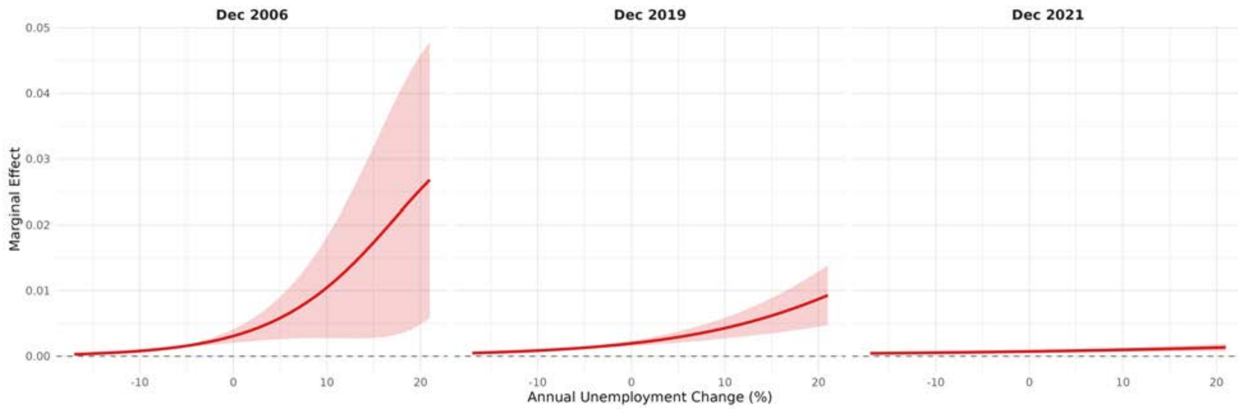

Figure 12: For a representative subprime borrower, the increase in the probability of transition from current to default for given year-over-year changes US state unemployment during December 2006, December 2019, and December 2021. Shaded areas represent 95 percent confidence intervals, where standard errors are clustered at the loan level. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

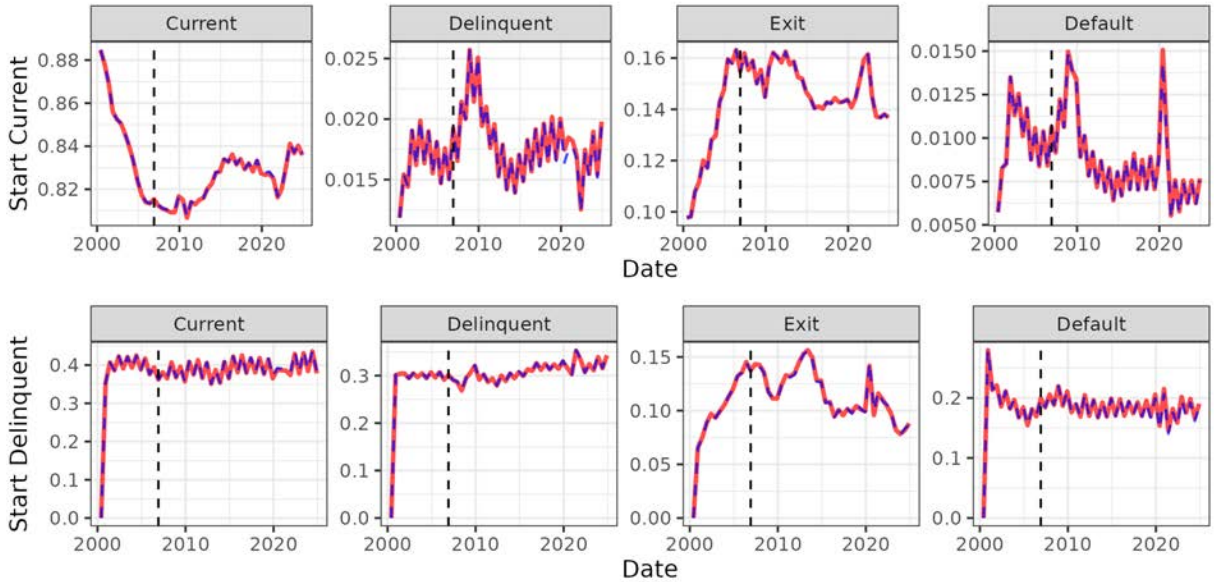

Figure 13: Estimated rolling regression transition probabilities on the full sample compared to the same estimation on on a sample that excludes loans that enter into forbearance. The red line represents pseudo out-of-sample predictions from the full sample and the blue dashed line represents pseudo out-of-sample predictions from a sample excluding loans in forbearance. Estimation variables include loan age, lender type, loan characteristics, and year-over-year changes in US state unemployment and HPI. Dashed line indicates December 2006, when rolling regressions start. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

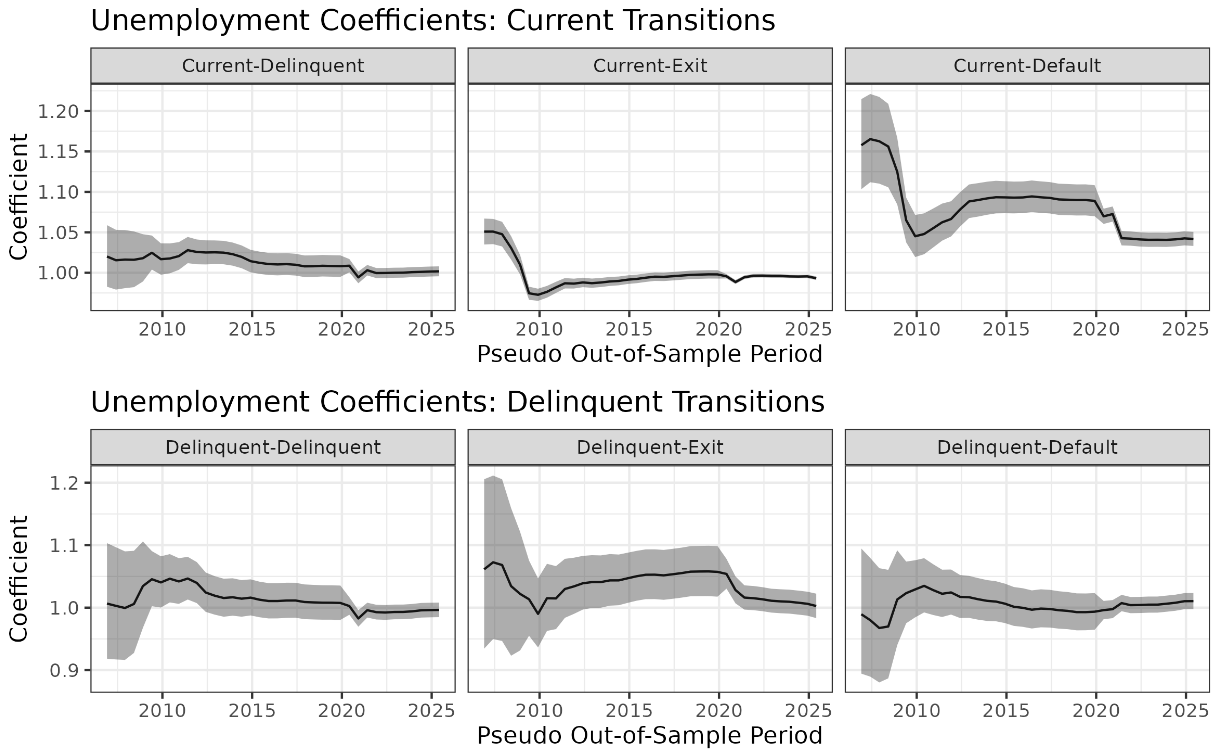

Figure 14: Covid Dummy Variable Results: Multinomial logit estimates of US state unemployment coefficients during pseudo out-of-sample estimations, 2012-2025. Estimation variables include loan age, lender type, loan characteristics, year-over-year changes in US state unemployment and HPI, and a dummy variable for 2000 and 2001. Shaded areas represent 95 percent confidence intervals, where standard errors are clustered at the loan level. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

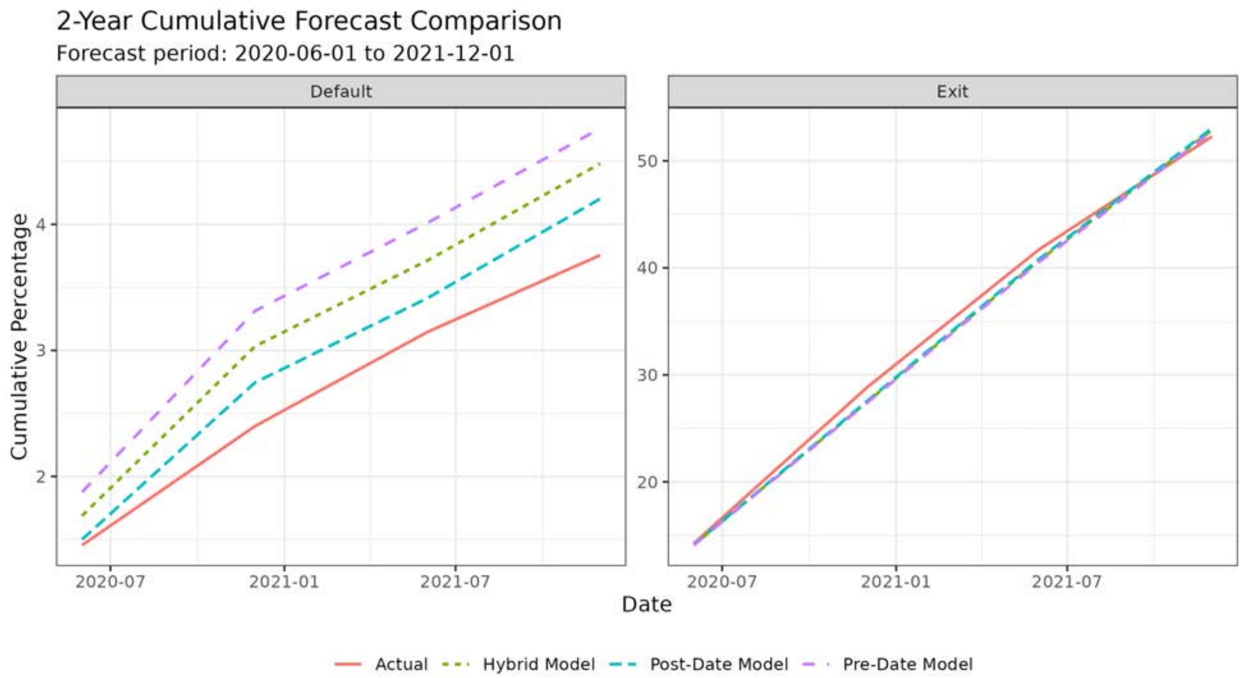

Figure 15: Cumulative 2-year forecasts from 2020-2021. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

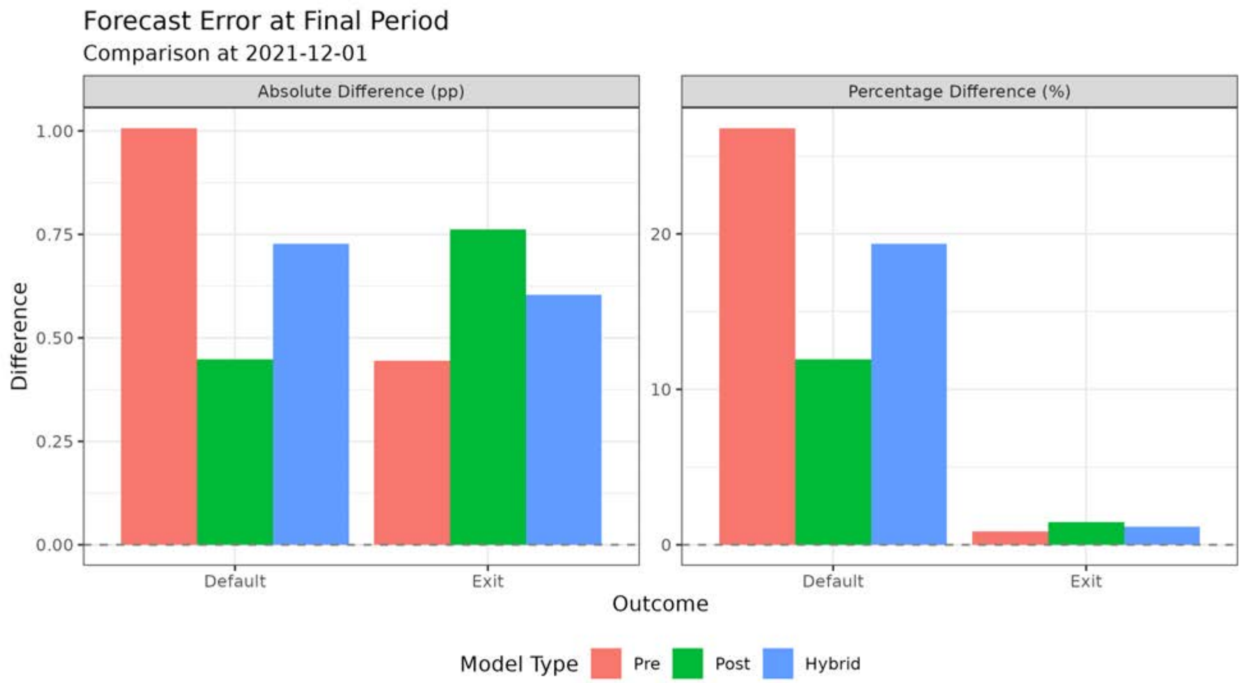

Figure 16: Cumulative 2-year forecast differences relative to actual outcomes. Cumulative forecasts over the period 2020-2021, differences computed as of December 2021. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

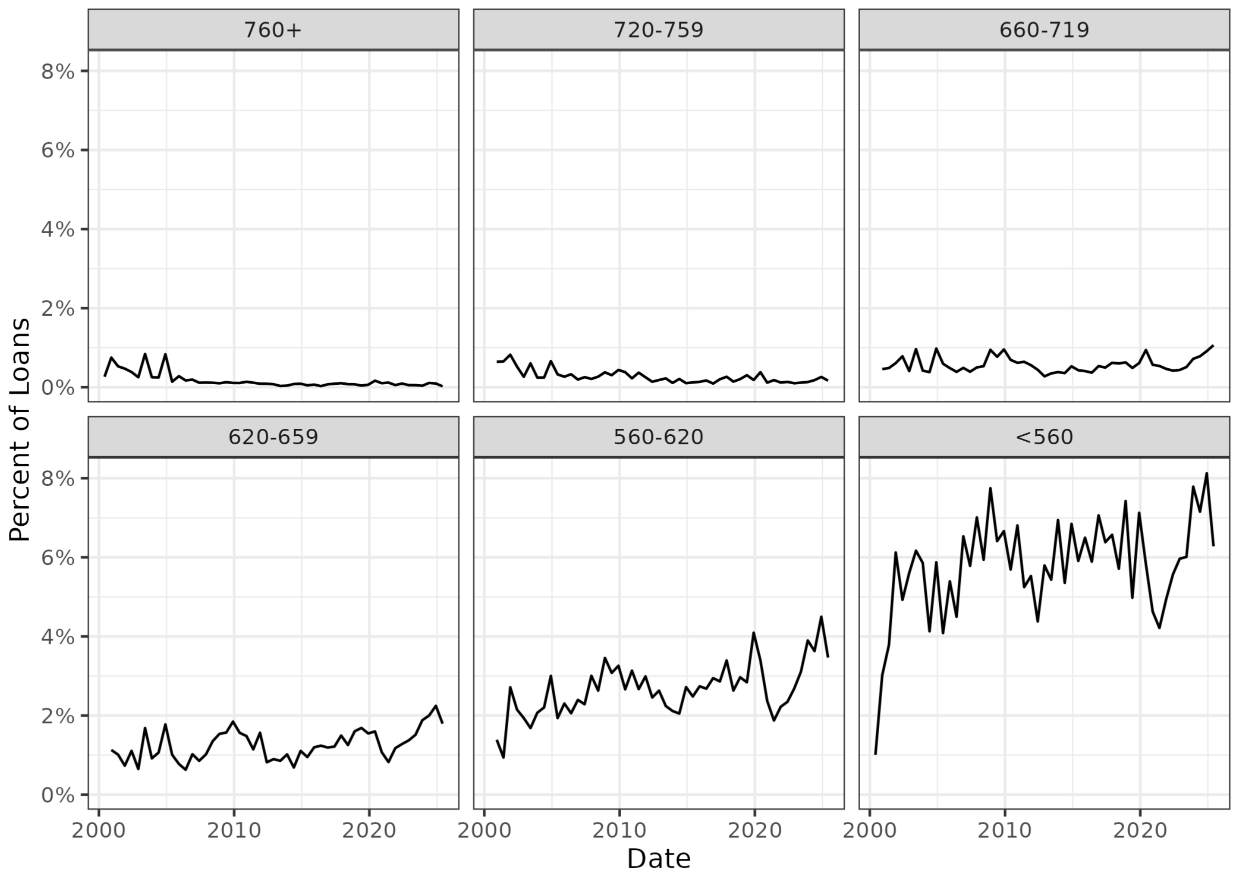

Figure 17: Default rates by original risk score group, December 2016-2025. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

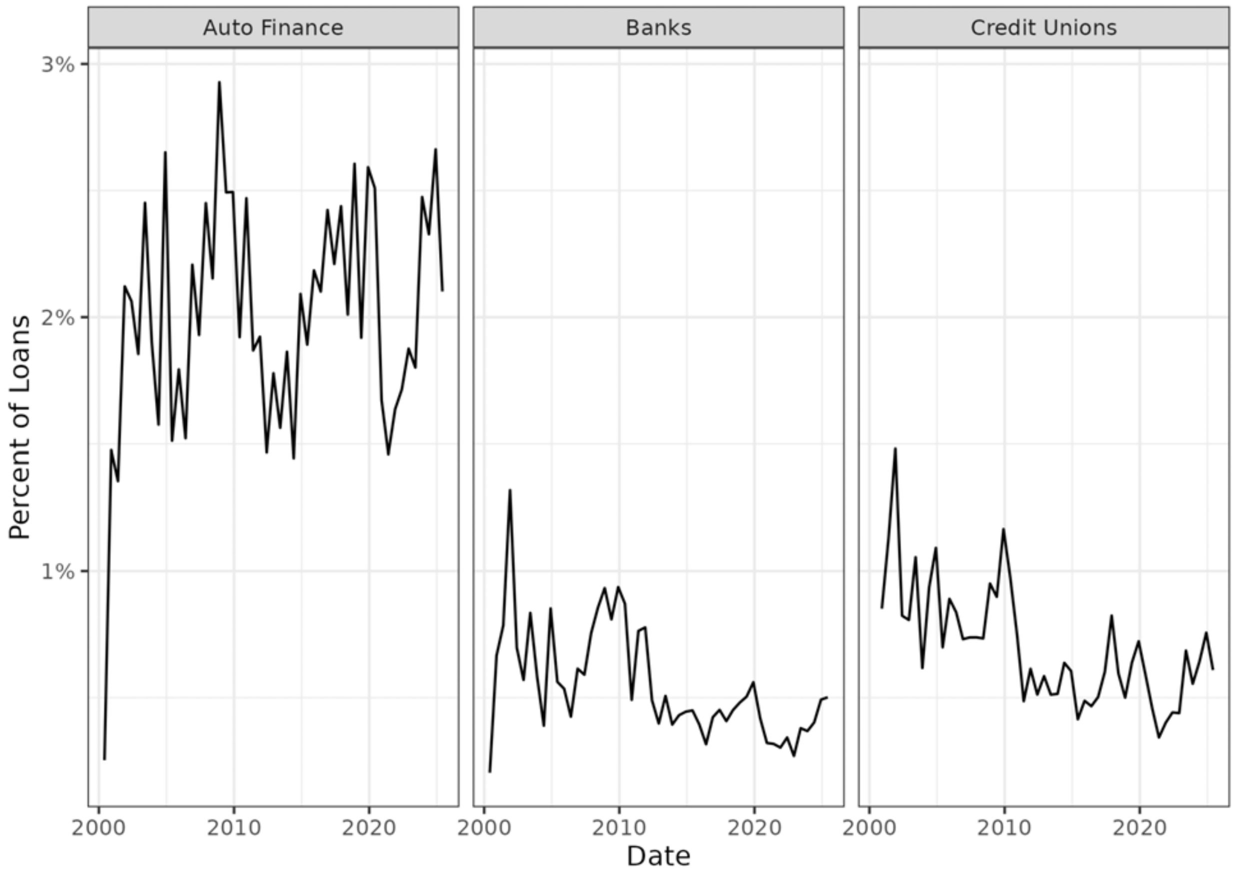

Figure 18: Default rates by lender type, December 2016-2025. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

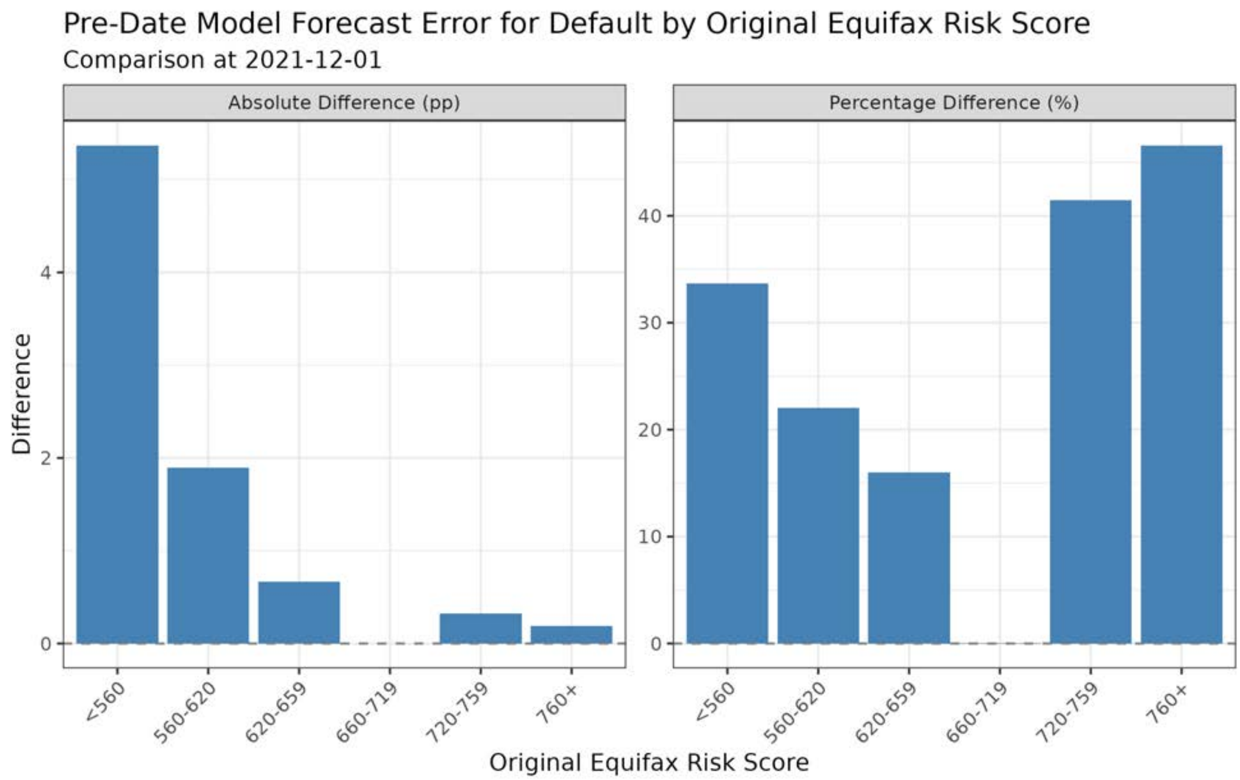

Figure 19: Cumulative 2-year forecast differences relative to actual outcomes, by original risk score groups. Cumulative forecasts over the period 2020-2021, differences computed as of December 2021. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

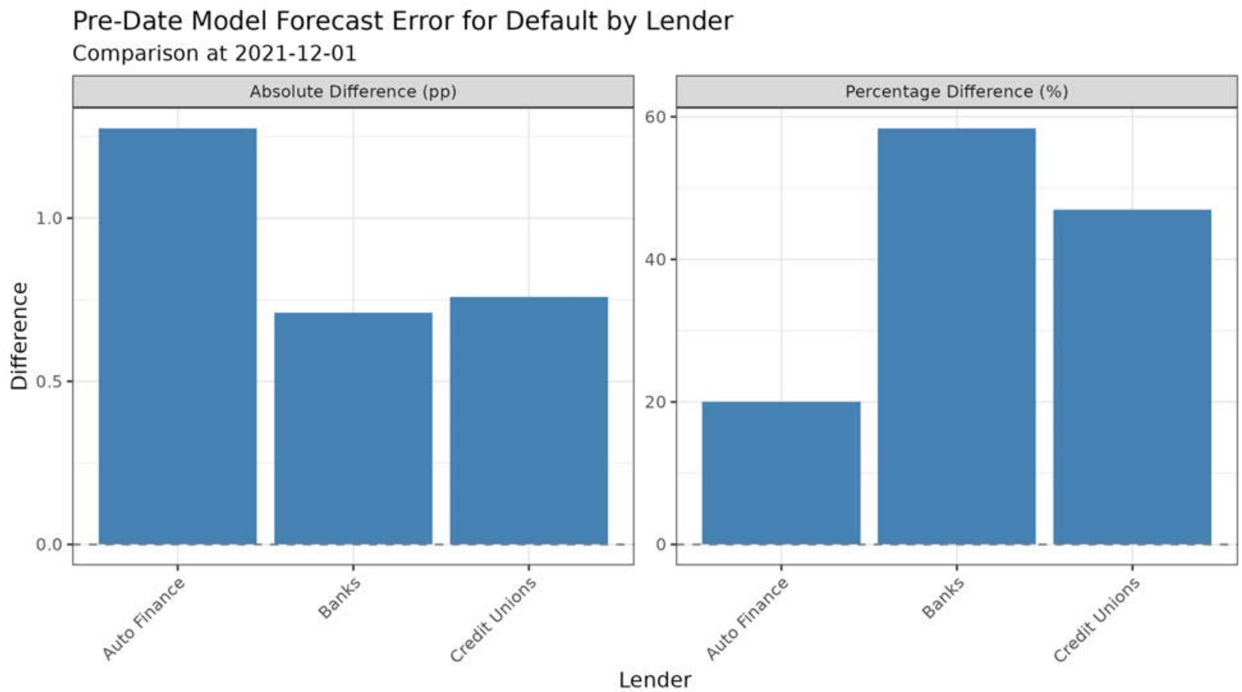

Figure 20: Cumulative 2-year forecast differences relative to actual outcomes, by lender. Cumulative forecasts over the period 2020-2021, differences computed as of December 2021. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

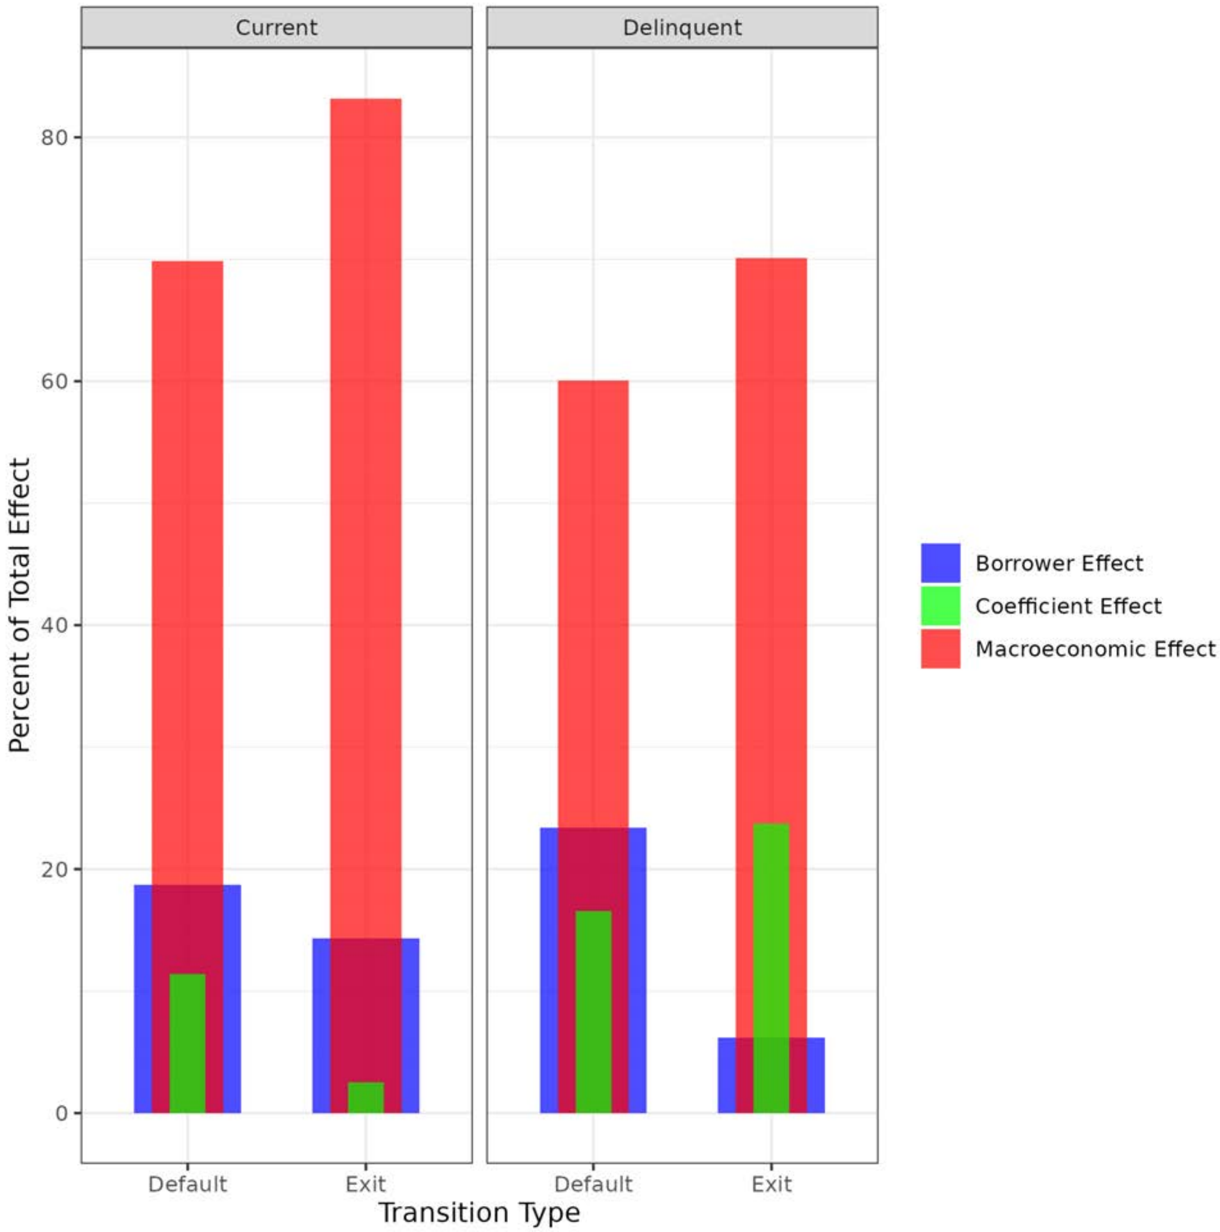

Figure 21: Oaxaca-Blinder decomposition on transition probabilities using the methods of Fairlie (2005). Baseline period is December 2019, alternative period is December 2021. Borrower and macro decompositions evaluated using alternative period coefficients. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

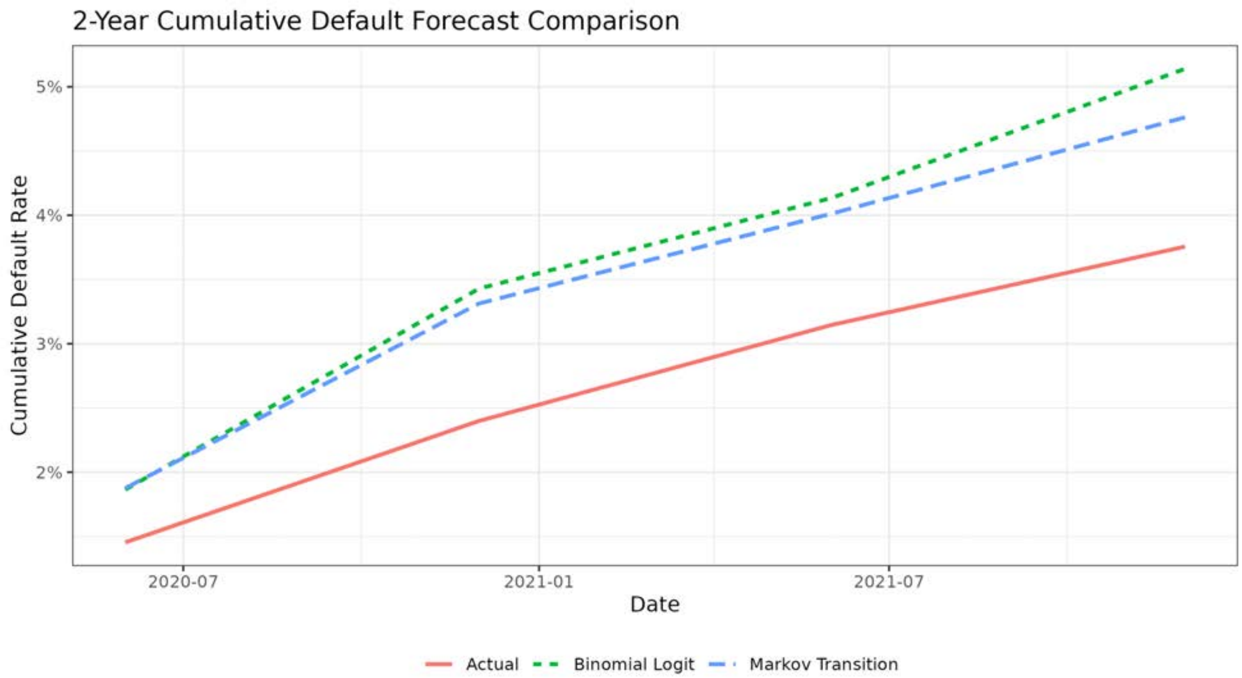

Figure 22: Comparison of Markov transition model and binomial logistic discrete-time hazard model cumulative 2-year forecasts from 2020-2021. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

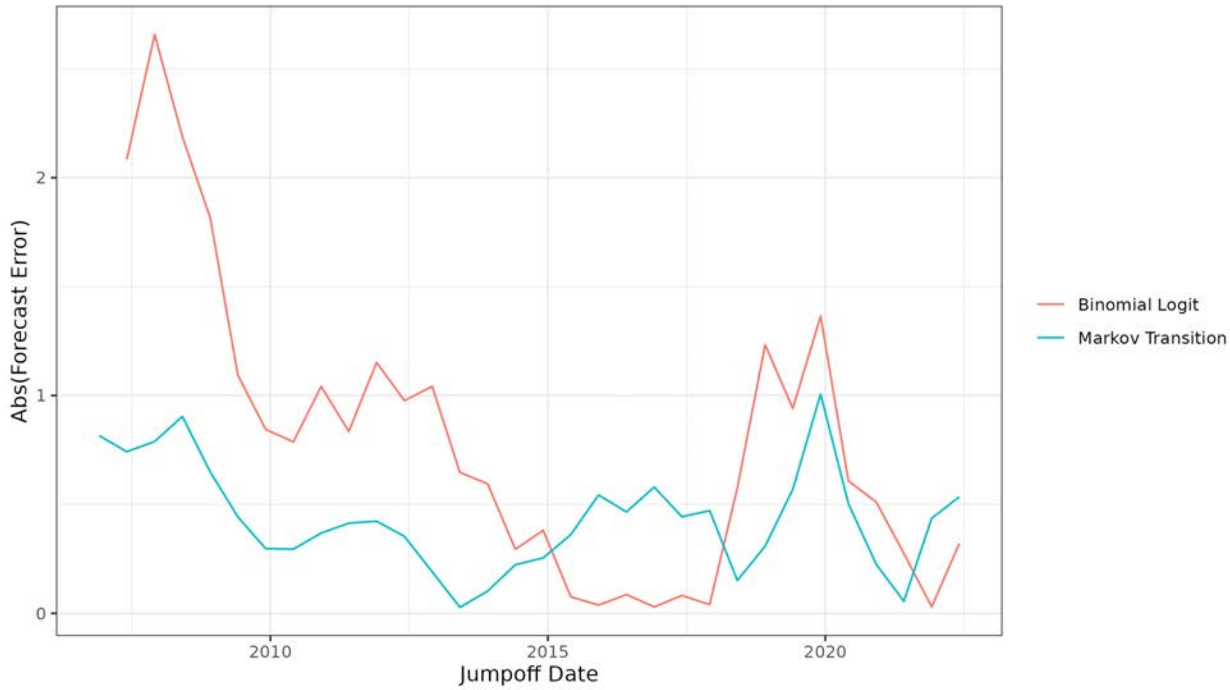

Figure 23: Comparison of rolling cumulative two-year default forecast errors between Markov transition model and binomial logistic discrete-time hazard model. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

A Appendix

A.1 Hazard Specification

Prior to estimating our full model, we examine how loan age may explain transition probabilities so that we can specify the functional form to use for age in the main model. Different choices in modeling the hazard function (such as exponential, Gompertz, Weibull or non-parametric) embed assumptions about how risk evolves over time. An incorrectly specified hazard can lead to biased estimates of covariate effects, misleading predictions about survival probabilities, and erroneous conclusions about duration dependence. For instance, imposing a monotonically increasing hazard function when the true hazard is non-monotonic can mask critical periods of heightened or reduced risk. As Lancaster (1990) demonstrates, these specification errors can fundamentally alter conclusions derived from the estimated model.

We begin by estimating a model that conditions transitions only on various transformations of loan age and excludes other covariates. We estimate the model over 2000-2019 to remove effects from the pandemic period. Following Wu et al. (2018), we compare the following specifications: Exponential, Gompertz, quadratic, Weibull, and nonparametric. The exponential form assumes a constant hazard; thus, the loan age drops out of the model altogether. The Gompertz assumes a log-linear relationship between the log odds of default and the hazard, and the quadratic offers more flexibility by allowing for non-monotonicity. The Weibull specification results in a log transformation to be placed on the loan age. Finally, we specify a nonparametric hazard by creating indicator variables for each year of the loan's life, grouping all loans greater than 5 years old into one final category due to a relatively small number of observations in that grouping.

We assess the in-sample fit of these competing specifications so that the data inform our choice of a hazard function for the main model. Our ideal specification will balance model fit and predictive accuracy with a parsimonious specification to avoid over-fitting. We evaluate model specification along these criteria using the root mean squared error (RMSE), McFadden's pseudo-R^2, and the Bayesian information criteria (BIC). Given that much of our focus is on aggregate outcomes, we compute the RMSE as the difference between the average predicted probabilities and actual relative frequency of outcomes. The pseudo-R^2 and BIC are computed normally using all of the observations.

Tables A1 and A2 show the model selection statistics for the current and delinquent models. Across both models the quadratic and nonparametric approaches generally have the best results. For the current state, the nonparametric specification has the smallest RMSE, while the quadratic has the largest pseudo-R^2 and smallest BIC. In the delinquent model, the quadratic has the lowest RMSE and BIC, and a pseudo-R^2 comparable to the nonparametric specification. Given these results, we choose the quadratic specification in our analysis for its simplicity and parsimonious specification relative to the nonparametric approach.

A.2 Regression Results: 2000-2006

To examine estimated relationships prior to the GFC, we produce tables of our multinomial logit regression results estimated over the 2000-2006 period. Tables A3 and A4 show the results of our estimation for transitions from current and delinquent, respectively. As before, the coefficients represent the marginal effects on the relative risk of transition to a particular state compared with staying in the current state. Credit unions and origination risk scores greater than 760 are the reference groups for lender type and origination risk score.

Although the magnitudes of various parameter estimates differ in magnitude compared to the 2000-2019 sample period, the estimated relationships for lender type and loan chacteristics that are statistically significant generally hold. For the macroeconomic variables, unemployment has a larger effect on the relative risk of transitioning from current-default in the 2000-2006 sample. Unemployment also has a counterintuitive effect on transitioning from both current-exit and delinquent-exit in the earlier sample, so that an increase in unemployment is associated with a higher relative risk of transition to that state. This could again be a result of measurement error from loans that were classified as exiting the sample but that actually defaulted. The estimates for HPI generally agree with the longer sample period except for the current-default state, where there is a small but positive relationshio between an increase in HPI and transitions from current-default. Across the macroeconomic variables, the absolute change in the unemployment coefficient for transitions from current-default between the 2000-2006 and 2000-2019 sample periods is relatively the largest.

A.3 Forecasts and Parameter Instability by Lender Type

We use the lender types associated with each loan in our data to examine if the relationship between unemployment and transitions to default and exit states changes differently in 2020 by lender type. To do this, we re-estimate the rolling out-of-sample framework on loans stratified by lender type so that we have lender-specific models and predictions. For statistical power, we focus on the broader lender types that we observe for our full sample period. Figures A1, A2, and A3 show these results for auto finance, bank, and credit union loans, respectively.

Focusing on transitions to default, loans originated from auto finance companies generally have more frequent occurrences of default relative to other lender types. The absolute forecast error for transitions from current to default is lowest for the auto finance loans, suggesting the relationship between unemployment and default transitions may have been more stable. For banks and credit unions, the absolute forecast errors for transitions from current to default in 2020 are larger, both more than 1 percentage point higher than for auto finance loans. The other notable source of forecast error in 2020 is transitions from delinquent to exit. This discrete jump in predicted exits during 2020 occurs primarily for auto finance companies.[14] We are cautious not to over-interpret those errors due to the potential for measurement error in that state.

Figures A4, A5, and A6 confirm the coefficient instability on the unemployment variable

[14] A discrete jump in the transitions from delinqeient to exit also occurs for credit union loans, but not until 2021 when the underlying HPI data began to spike.

is associated with the size of the forecast errors by lender. The largest amount of instability for transitions from current to default is observed for banks and credit unions. In both cases, the marginal effect of the year-over-year change in unemployment decreases by roughly 15 percentage points in 2020. After 2020, the estimated effect of unemployment is statistically zero for banks and credit unions, and remains that way until the end of the sample period. The rolling coefficients for unemployment on delinquent to exit transitions are statistically significant only for auto finance loans prior to 2020, and drop by roughly 10 percentage points to statistically zero levels after 2020. Despite some volatility in 2020, these coefficients are statistically zero for other lender types for most of the out-of-sample period.

A.4 Forecasts and Parameter Instability with Alternative Macroeconomic Specifications

A.4.1 Model without macroeconomic variables

Given the instability in the unemployment coefficients and the resulting forecast errors, it is instructive to understand how a model that excludes these variables might perform. We re-estimate the rolling regression model excluding macroeconomic variables, again treating June 2000-December 2006 as the initial in-sample period.

Figure A7 plots the resulting semiannual out-of-sample forecasts. Not surprisingly, forecast errors for transition to default decline during the GFC and Covid recessions. However, forecast errors worsen for transitions to the exit state. This is reinforced by the two-year cumulative forecast from 2019-2021 using this model, as shown in Figure A8. Relative to the model that includes macroeconomic variables, the cumulative forecast errors are smaller across models that use parameters from before or after 2020, but all of the parameter choices perform worse for the exit state. This suggests that a model without macro variables may improve forecasts for certain absorbing state transitions, but at the expense of others.

A.4.2 Model with additional macroeconomic variables

As an alternative to excluding macroeconomic variables, we consider additional macroeconomic variables beyond just year-over-year changes in unemployment and HPI to examine if the additional macroeconomic dynamics improve forecast accuracy. Following Heitfield and Sabarwal (2004), we include the annual difference in one-year Treasury rates to capture incentives to refinance during periods of low interest rates. While our data do not include information on loan-to-value ratios, we attempt to capture changes in the market value of auto loans by including the annual percentage change in the consumer price index for used vehicles obtained from the Bureau of Labor Statistics.[15] We also control for the level of unemployment as opposed to only the year-over-year change in order to capture both observable deterioration and severity of unemployment at a given time. Figure A9 plots the resulting semiannual out-of-sample forecasts. The forecasts are nearly identical to the forecasts that include only unemployment and HPI. This offers some evidence that the forecast errors and parameter instability are not a result of additional macroeconomic dynamics missing from the model.

[15] Specifically, we use the consumer price index for all urban consumers based on average prices of used cars and trucks in US cities.

Table A1: Current State Model Hazard Fit Comparison.

Model	RMSE	Pseudo-R2	BIC
Exponential	17.3	0	1292080
Gompertz	6.2	0.048	1230154
Quadratic	7.4	0.051	1226555
Weibull	11.7	0.034	1248651
Nonparametric	6	0.048	1230430

Fit statistics for various specifications of the hazard function in the model of transitions from the current state. All specifications estimated only with a hazard function and no other covariates to determine best fit.

Table A2: Delinquent State Model Hazard Fit Comparison.

Model	RMSE	Pseudo-R2	BIC
Exponential	12.5	0	76895
Gompertz	7.5	0.021	75304
Quadratic	6.5	0.025	75062
Weibull	8.8	0.019	75448
Nonparametric	7.1	0.025	75188

Fit statistics for various specifications of the hazard function in the model of transitions from the delinquent state. All specifications estimated only with a hazard function and no other covariates to determine best fit.

Table A3: Current State Multinomial Logit Results (2000-2006)

	y = Transitions from Current State		
	Current-Delinquent	Current-Exit	Current-Default
Intercept	0.000***	74.180***	0.044***
	[0.000, 0.000]	[73.993, 74.368]	[0.044, 0.044]
Lender: Auto Finance	1.487***	0.905***	1.290***
	[1.425, 1.551]	[0.880, 0.931]	[1.253, 1.329]
Lender: Banks	1.098***	0.951***	0.748***
	[1.061, 1.136]	[0.923, 0.981]	[0.733, 0.763]
Loan Age (months)	1.057***	1.006***	1.083***
	[1.050, 1.064]	[1.003, 1.008]	[1.073, 1.093]
Loan Age ^{2}	0.999***	1.000***	0.999***
	[0.999, 1.000]	[1.000, 1.000]	[0.999, 0.999]
Equifax Risk Score: 720-759	2.377***	1.015	0.943***
	[2.371, 2.383]	[0.981, 1.050]	[0.939, 0.947]
Equifax Risk Score: 660-719	6.693***	1.022	1.253***
	[6.602, 6.784]	[0.990, 1.055]	[1.242, 1.265]
Equifax Risk Score: 620-659	16.983***	1.054***	1.975***
	[16.323, 17.670]	[1.016, 1.095]	[1.960, 1.989]
Equifax Risk Score: <560	84.648***	0.943***	8.319***
	[81.225, 88.215]	[0.904, 0.984]	[7.970, 8.683]
Borrower Age	0.993***	0.996***	0.991***
	[0.990, 0.995]	[0.995, 0.997]	[0.988, 0.994]
First Mortgage Flag (t-1)	0.891***	1.079***	0.739***
	[0.860, 0.924]	[1.054, 1.105]	[0.731, 0.746]
Ln(Balance) (t-1)	1.140***	0.512***	0.763***
	[1.125, 1.155]	[0.509, 0.515]	[0.750, 0.776]
Joint Account	0.889***	0.939***	0.870***
	[0.842, 0.940]	[0.919, 0.960]	[0.864, 0.876]
Annual Unemployment Change	1.020	1.051***	1.158***
	[0.983, 1.059]	[1.035, 1.067]	[1.105, 1.213]
Annual HPI Change	0.987***	1.009***	0.993**
	[0.982, 0.993]	[1.007, 1.011]	[0.986, 1.000]
Seasonal Effect (Jan-June = 1)	0.828***	1.044***	0.814***
	[0.783, 0.877]	[1.022, 1.066]	[0.810, 0.818]
Num.Obs.	305003		
R2	0.081		
R2 Adj.	0.081		
Pseudo-R2	0.08		
AUC	0.69		

Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations. Model estimated on sample data from 2000-2006. Coefficients represent relative risks. *, **, *** represent statistical significance at 0.1, 0.05, and 0.01 levels, respectively. 95 percent confidence intervals shown in brackets. Standard errors are clustered at the tradeline level.

Table A4: Delinquent State Multinomial Logit Results (2000-2006)

	y = Transitions from Delinquent State		
	Delinquent-Delinquent	Delinquent-Exit	Delinquent-Default
Intercept	0.277***	17.602***	19.453***
	[0.276, 0.278]	[17.509, 17.695]	[19.349, 19.556]
Lender: Auto Finance	1.221***	0.847***	0.938
	[1.104, 1.349]	[0.753, 0.954]	[0.829, 1.061]
Lender: Banks	1.010	1.115**	0.768***
	[0.933, 1.094]	[1.004, 1.239]	[0.701, 0.842]
Loan Age (months)	0.982*	0.985	0.951***
	[0.963, 1.002]	[0.960, 1.011]	[0.930, 0.973]
Loan Age ^{2}	1.000***	1.001***	1.001***
	[1.000, 1.001]	[1.000, 1.001]	[1.000, 1.001]
Equifax Risk Score: 720-759	1.428***	0.522***	2.044***
	[1.414, 1.442]	[0.518, 0.526]	[2.023, 2.066]
Equifax Risk Score: 660-719	1.151*	0.535***	1.428***
	[0.983, 1.349]	[0.452, 0.633]	[1.235, 1.651]
Equifax Risk Score: 620-659	1.759***	0.652***	1.695***
	[1.526, 2.028]	[0.546, 0.779]	[1.427, 2.014]
Equifax Risk Score: <560	2.306***	0.683***	3.173***
	[2.075, 2.564]	[0.594, 0.785]	[2.798, 3.598]
Borrower Age	1.001	0.999	0.989***
	[0.996, 1.006]	[0.992, 1.006]	[0.982, 0.995]
First Mortgage Flag (t-1)	0.850**	0.952	0.647***
	[0.726, 0.996]	[0.769, 1.178]	[0.526, 0.796]
Ln(Balance) (t-1)	1.039*	0.671***	0.726***
	[0.998, 1.081]	[0.635, 0.709]	[0.694, 0.759]
Joint Account	0.932	0.953	0.780***
	[0.820, 1.059]	[0.799, 1.137]	[0.669, 0.910]
Annual Unemployment Change	1.007	1.061	0.989
	[0.922, 1.099]	[0.939, 1.200]	[0.893, 1.096]
Annual HPI Change	0.996	0.994	0.981***
	[0.985, 1.008]	[0.979, 1.010]	[0.967, 0.994]
Seasonal Effect (Jan-June = 1)	0.886*	0.941	0.801***
	[0.783, 1.003]	[0.793, 1.117]	[0.693, 0.927]
Num.Obs.	5992		
R2	0.030		
R2 Adj.	0.030		
Pseudo-R2	0.03		
AUC	0.61		

Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations. Model estimated on sample data from 2000-2006. Coefficients represent relative risks. *, **, *** represent statistical significance at 0.1, 0.05, and 0.01 levels, respectively. 95 percent confidence intervals shown in brackets. Standard errors are clustered at the tradeline level.

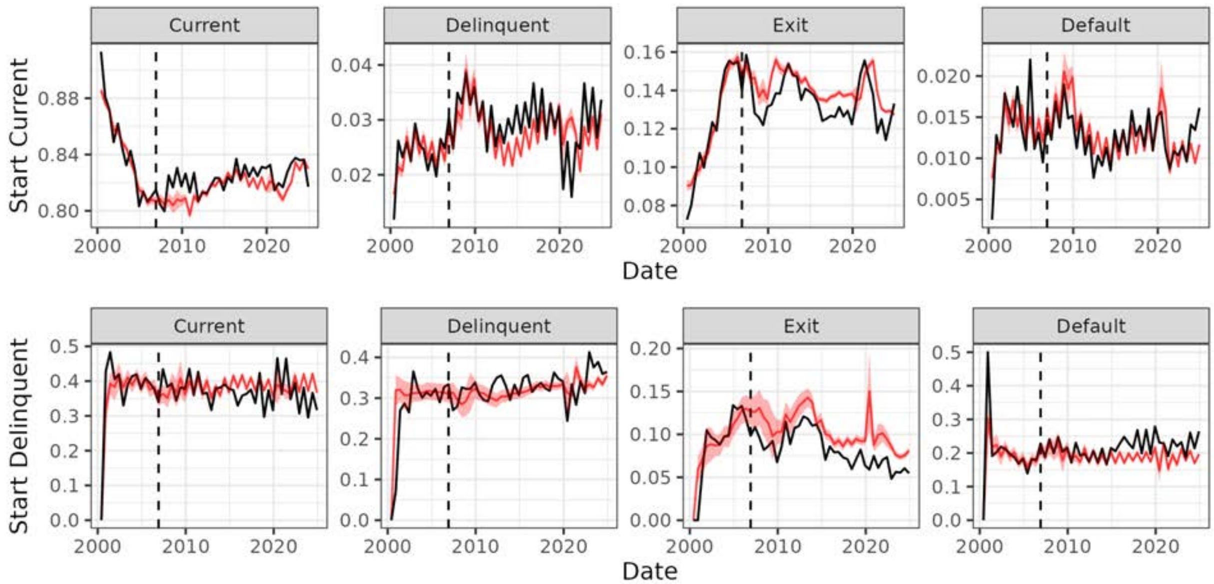

Figure A1: Auto finance loans: Time path of current and delinquent-state transition probabilities from pseudo out-of-sample estimation. Estimation variables include loan age, lender type, loan characteristics, and year-over-year changes in US state unemployment and HPI. The red line represents pseudo out-of-sample forecasts and the black line represents actual outcomes. Red shaded areas represent 95 percent confidence intervals, where standard errors are clustered at the loan level. Dashed line indicates final in-sample forecast as of December 2006. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

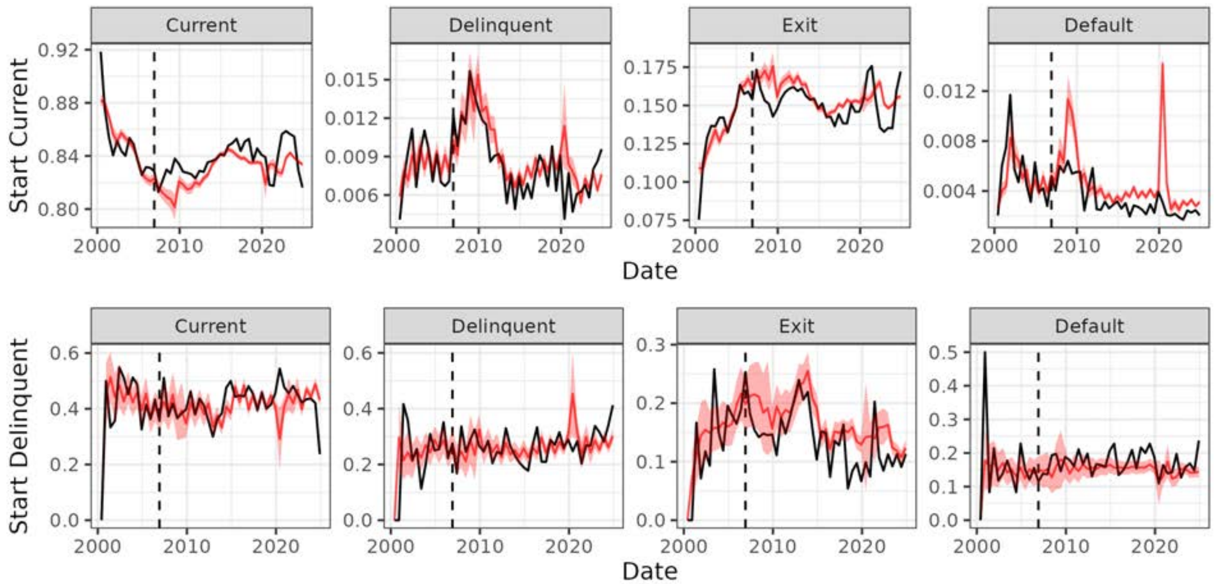

Figure A2: Bank loans: Time path of current and delinquent-state transition probabilities from pseudo out-of-sample estimation. Estimation variables include loan age, lender type, loan characteristics, and year-over-year changes in US state unemployment and HPI. The red line represents pseudo out-of-sample forecasts and the black line represents actual outcomes. Red shaded areas represent 95 percent confidence intervals, where standard errors are clustered at the loan level. Dashed line indicates final in-sample forecast as of December 2006. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

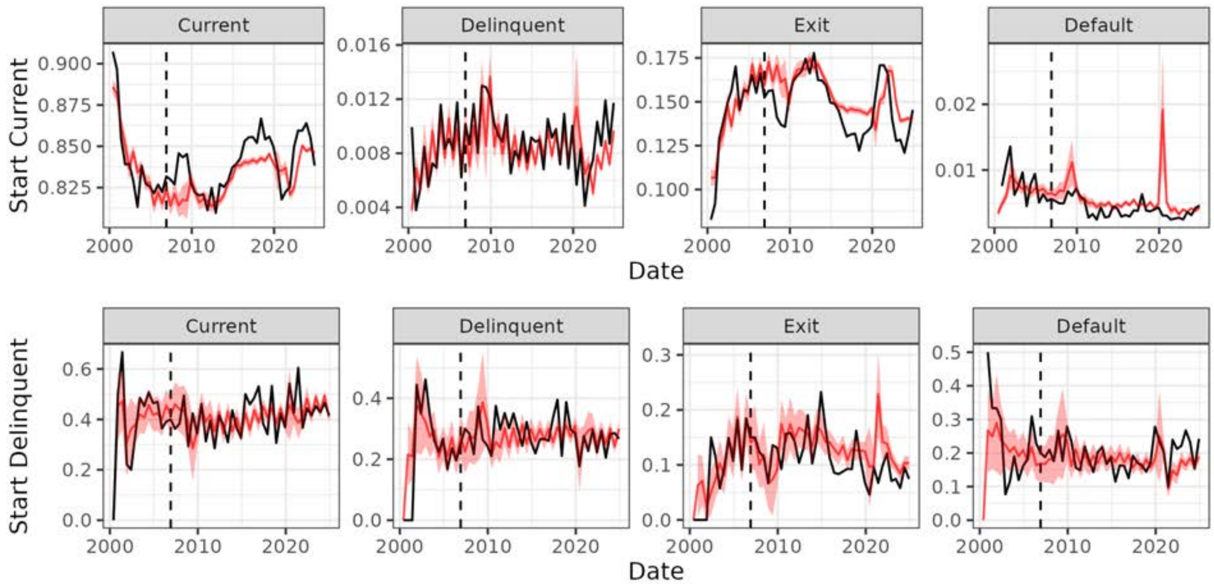

Figure A3: Credit union loans: Time path of current and delinquent-state transition probabilities from pseudo out-of-sample estimation. Estimation variables include loan age, lender type, loan characteristics, and year-over-year changes in US state unemployment and HPI. The red line represents pseudo out-of-sample forecasts and the black line represents actual outcomes. Red shaded areas represent 95 percent confidence intervals, where standard errors are clustered at the loan level. Dashed line indicates final in-sample forecast as of December 2006. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

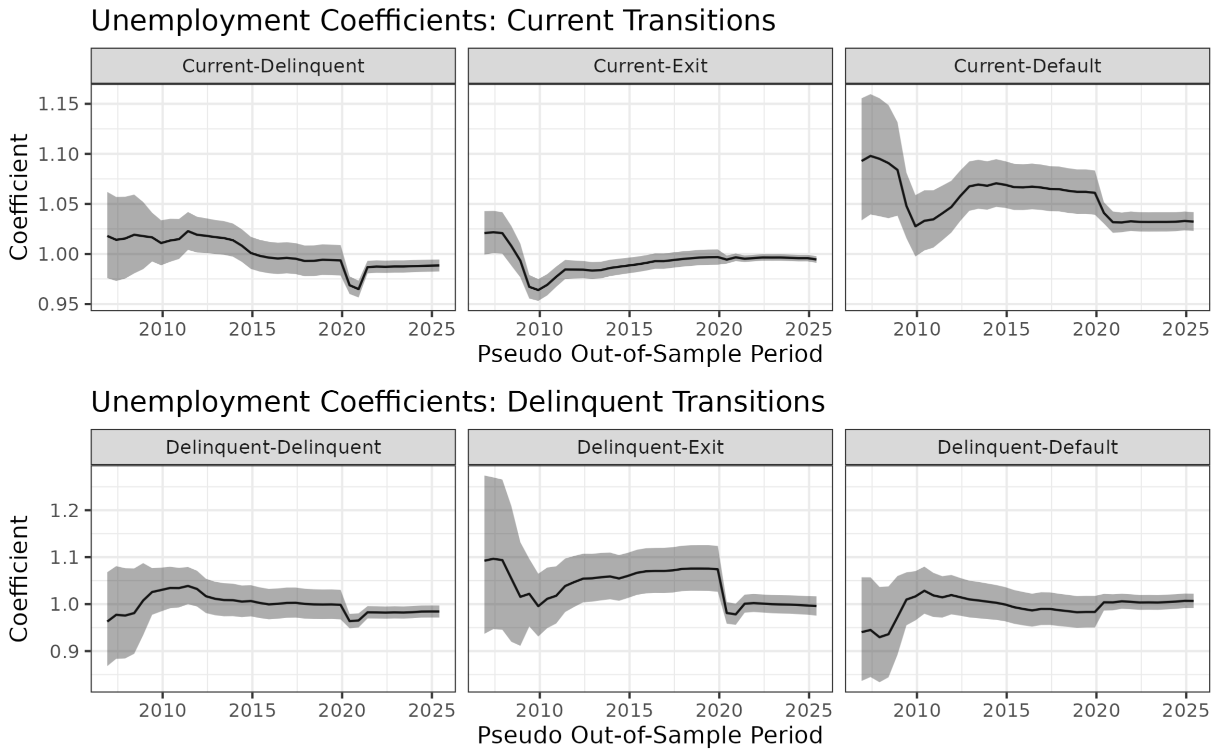

Figure A4: Auto finance loans: Multinomial logit estimates of US state unemployment coefficients during pseudo out-of-sample estimations, 2006-2025. Estimation variables include loan age, lender type, loan characteristics, and year-over-year changes in US state unemployment and HPI. Shaded areas represent 95 percent confidence intervals, where standard errors are clustered at the loan level. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

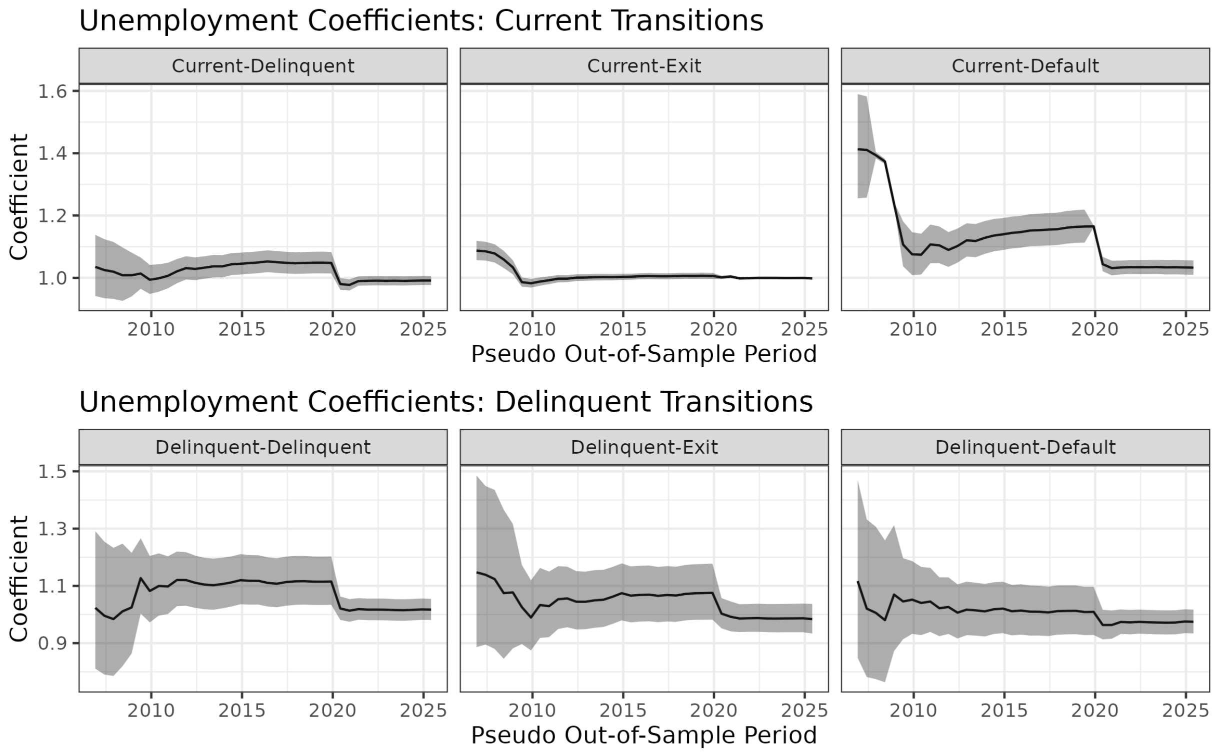

Figure A5: Bank loans: Multinomial logit estimates of US state unemployment coefficients during pseudo out-of-sample estimations by model, 2006-2025. Estimation variables include loan age, lender type, loan characteristics, and year-over-year changes in US state unemployment and HPI. Shaded areas represent 95 percent confidence intervals, where standard errors are clustered at the loan level. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

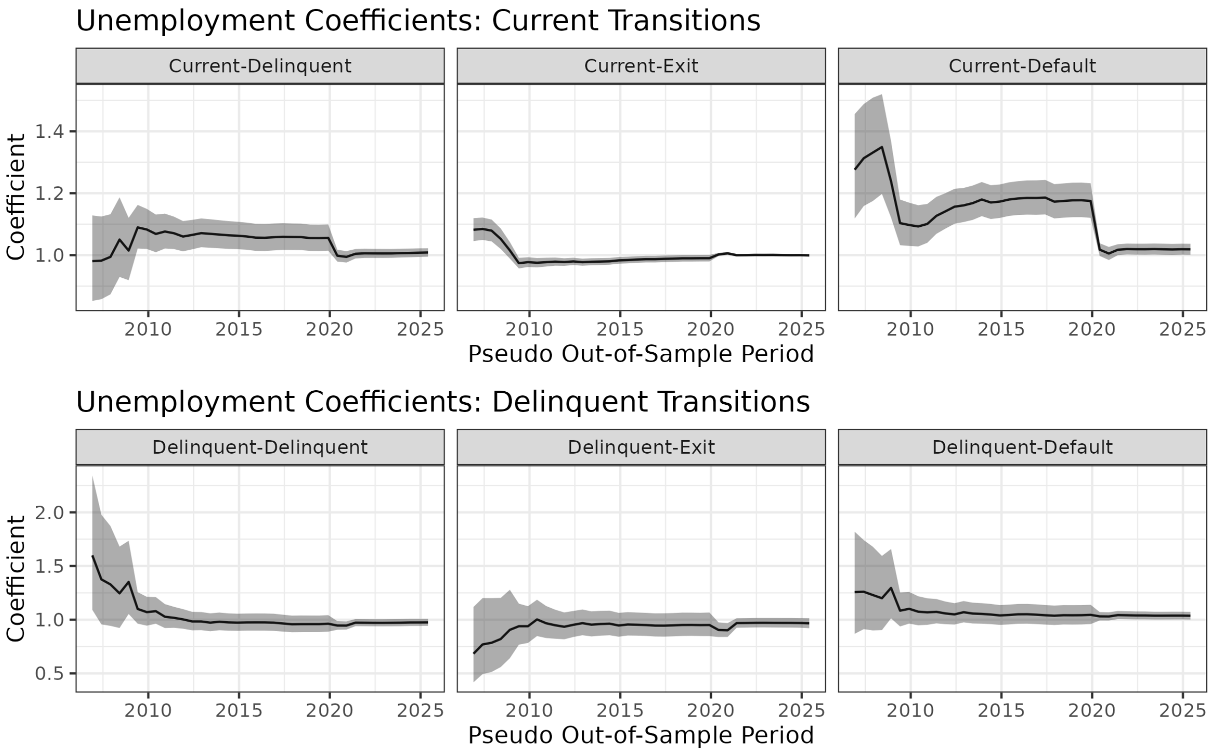

Figure A6: Credit union loans: Multinomial logit estimates of US state unemployment coefficients during pseudo out-of-sample estimations, 2006-2025. Estimation variables include loan age, lender type, loan characteristics, and year-over-year changes in US state unemployment and HPI. Shaded areas represent 95 percent confidence intervals, where standard errors are clustered at the loan level. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

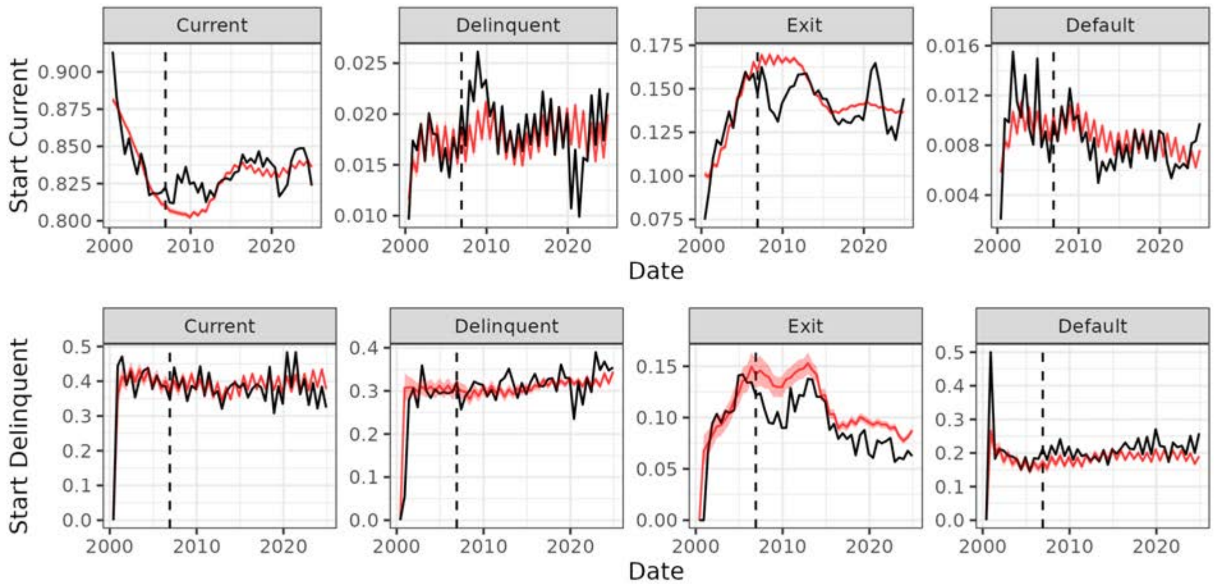

Figure A7: Model without macroeconomic variables: Time path of current and delinquent-state transition probabilities from pseudo out-of-sample estimation. Estimation variables include loan age, lender type, loan characteristics, but excludes any macroeconomic variables. The red line represents pseudo out-of-sample forecasts and the black line represents actual outcomes. Red shaded areas represent 95 percent confidence intervals, where standard errors are clustered at the loan level. Dashed line indicates final in-sample forecast as of December 2006. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

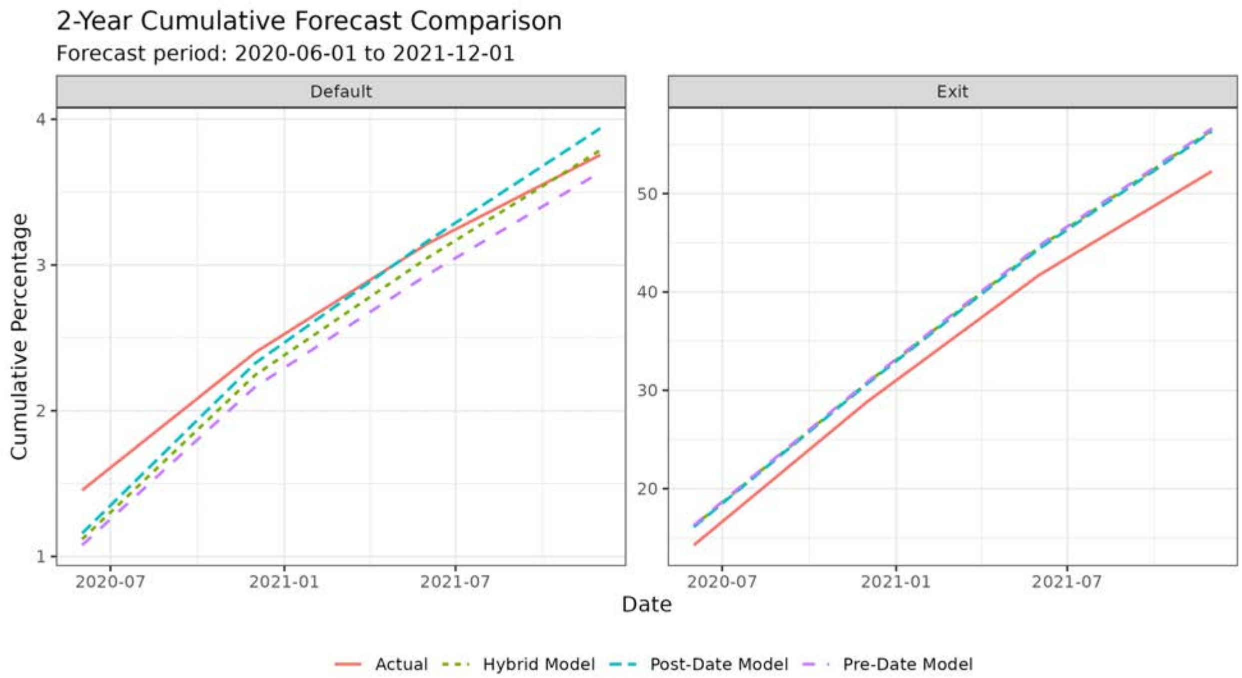

Figure A8: Model without macroeconomic variables: Cumulative 2-year forecast differences relative to actual outcomes. Cumulative forecasts over the period 2020-2021, differences computed as of December 2021. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations

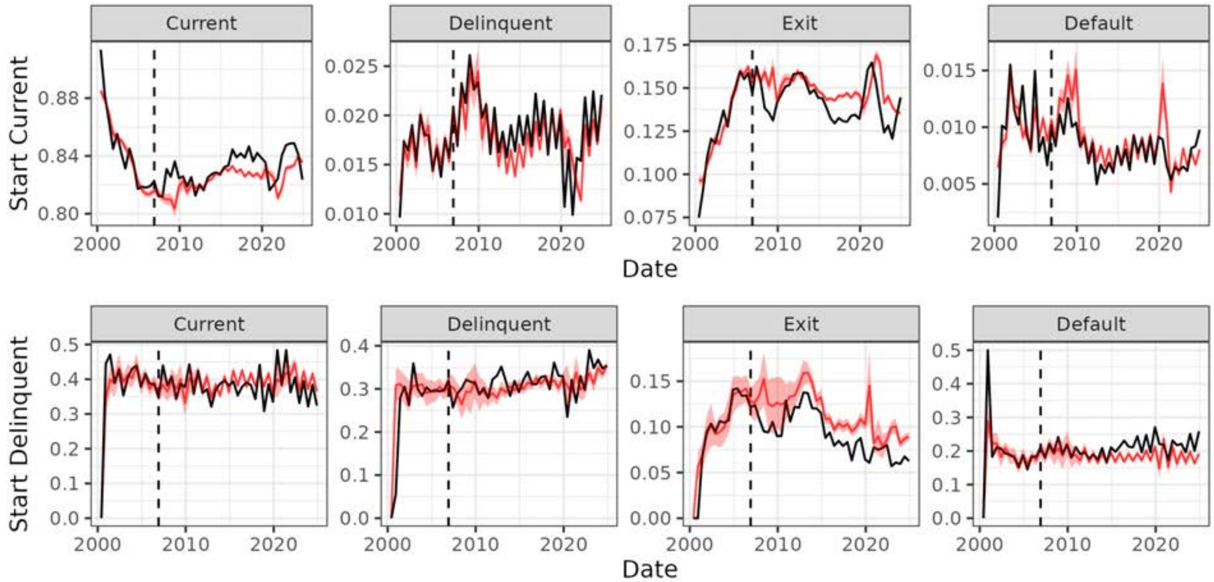

Figure A9: Model without additional variables: Time path of current and delinquent-state transition probabilities from pseudo out-of-sample estimation. Loan-level estimation variables include loan age, lender type, loan characteristics. Macroeconomic variables include the level of unemployment, the absolute year-over-year change in US state unemployment, the 1-year Treasury rate, the annual percentage change in HPI, and the consumer price index for used vehicles. The red line represents pseudo out-of-sample forecasts and the black line represents actual outcomes. Red shaded areas represent 95 percent confidence intervals, where standard errors are clustered at the loan level. Dashed line indicates final in-sample forecast as of December 2006. Source: FRBNY Consumer Credit Panel / Equifax Data, Authors' calculations