When Models Fail: Evidence from Automated Underwriting in Auto Loan Markets

July 1, 2025

Abstract

While prior studies find the outperformance of automated over manual underwriting, I document significant heterogeneity in the adoption of automated underwriting practices across auto lenders and within the same lender. To explain this heterogeneity, I examine the performance of automated underwriting systems under conditions of heightened data uncertainty caused by the COVID-19 pandemic. Using a difference-in-differences design, I estimate the impact of this unprecedented shock on the performance of automated underwriting in the auto loan market. My findings show that the performance of automated underwriting, as measured by ex-post default rates, deteriorated substantially relative to human underwriters during this period. This effect is particularly pronounced among low-income borrowers. These results suggest the limitations of automated underwriting systems when faced with unprecedented shocks outside the scope of their historical training datasets, underscoring the continued relevance of human underwriters in addressing such challenges in the auto lending industry.

1 Introduction

The integration of advanced technologies in consumer credit markets has significantly reshaped traditional lending practices. Recent studies show that automated underwriting enhances operational efficiency by enabling faster processing of applications without compromising default risk (Fuster et al. 2019), reduces discriminatory practices in credit decisions (Howell et al. 2024), and promotes financial inclusion by extending credit access to high-risk borrowers without increasing default probabilities (Gao, Yi, and Zhang 2024). Furthermore, automation contributes to improved lender profitability by mitigating agency conflicts and providing higher capacity to process more complex credit applications (Jansen, Nguyen, and Shams 2024).

Given these advantages, it would be reasonable to expect lenders to fully utilize automated systems in their operations to capitalize on their potential benefits. However, the extent of automation adoption varies considerably across lenders. While some lenders heavily rely on automated systems to process a significant proportion of their credit applications, others continue to depend predominantly on human decision-making in loan origination. This heterogeneity presents an intriguing puzzle and raises critical questions: What drives some lenders to refrain from fully automating their underwriting practices? Are these institutions failing to realize potential economic gains by not capitalizing on the efficiencies and benefits of automation?

Understanding why lenders do not fully automate their underwriting processes is crucial and necessitates further analysis. The main hypothesis of this study is that the observed heterogeneity in the adoption of automation stems from the limitations of automated systems during periods of economic uncertainty. While a growing body of literature has examined the role of technology in household consumer markets, most studies focus on its application during stable periods, leaving the comparative performance of technology during large-scale shocks relatively unexplored. In this paper, I investigate the performance of automated versus human underwriting in the context of large-scale shocks, offering insights into how uncertainty shapes the relative effectiveness of these systems. Automated models, trained on historical data to assess borrower creditworthiness, often struggle to adapt to the rapidly changing conditions induced by large-scale shocks. When historical data becomes less relevant or fails to accurately reflect current circumstances, the predictive accuracy of these systems deteriorates. Furthermore, automated systems require time to recalibrate, as updating datasets, retraining models, and integrating new information is a time-intensive process. By contrast, human underwriters possess the ability to interpret and adapt to new and evolving information in real time, enabling more responsive and effective decision-making during periods of uncertainty.

To test my hypothesis, I use the U.S. auto loan market as a natural laboratory and a difference-in-differences design to evaluate the performance of loans originated by automated systems and human underwriters before and after the COVID-19 pandemic, a significant large-scale shock. The U.S. auto loan market, the second-largest segment of consumer credit with $1.64 trillion in outstanding debt as of Q3 2024 (FRBNY), provides an ideal setting for this analysis due to its scale and broad implications for both lenders and consumers. This market uniquely features both human and automated underwriting practices, enabling a controlled comparison of loan performance within the same lender, thereby eliminating borrower self-selection bias—a limitation of studies comparing FinTech and traditional banks. Additionally, the Regulation AB-II dataset, which mandates detailed loan-level reporting for securitized auto loans, allows precise identification of underwriting methods. The COVID-19 pandemic offers a valuable setting to assess these systems under economic instability, marking the first major, unexpected systemic shock faced by automated systems that had previously evolved under stable conditions. Furthermore, the Regulation AB-II dataset, which covers loans originated after 2017, aligns with this timeline, enabling an analysis of performance during both stable and unstable economic conditions.

Assessing the relative performance of automated and human underwriting before and after the COVID-19 shock poses a key empirical challenge, as lenders typically assign these underwriting methods to distinct borrower profiles. In an ideal setting, applicants would be randomly assigned to human or automated underwriting both before and after the pandemic, enabling a difference-in-differences design to capture shifts in their relative performance under economic stress. In the absence of such an experiment, I use lender-specific discontinuities in the likelihood of automation across various FICO scores as a source of quasi-exogenous variation. Given that borrowers' observable characteristics are smooth across these thresholds, the discontinuities generate quasi-random variation in the likelihood of automation among otherwise comparable borrowers, allowing me to isolate the causal effect of automation on loan performance. To identify these discontinuities, I employ a data-driven procedure and utilize their locations as instrumental variables within a fuzzy regression discontinuity design.

I then incorporate this fuzzy RDD into a difference-in-differences framework to examine how the causal effect of automation varies across economic conditions. Fundamentally, this approach leverages the comparison between loans originated just above the FICO score thresholds—where automated underwriting is predominantly employed—and those originated just below the thresholds—where human underwriting is more commonly applied—across the pre- and post-pandemic periods. Consistent with prior research (Jansen, Nguyen, and Shams 2024), I find that automated underwriting outperforms human underwriting during stable economic periods, even among comparable borrowers within the same lender. However, this dynamic shifts dramatically after the onset of the COVID-19 pandemic. My findings reveal that the performance advantage of automated systems not only diminishes but reverses, with human underwriters outperforming automated systems in the post-pandemic period. This reversal likely reflects the limitations of automated systems, which rely on models trained on historical data and operate under the assumption of stable economic conditions. These systems appear to have struggled to adapt to the abrupt and unprecedented disruptions caused by the pandemic. In contrast, human underwriters, with their ability to incorporate real-time information and exercise discretion in assessing risk, might have been more adaptable to the changing conditions.

Why does the relative performance of automated underwriting systems deteriorate after the onset of the COVID-19 pandemic? To explore this question, I start by examining the role of borrower income through cross-sectional analysis. Lower-income borrowers, who are inherently more susceptible to economic disruptions, faced significant challenges during the pandemic, including widespread job losses and reductions in earnings. These factors likely increased their financial vulnerability and default risk. An analysis of default rates across income deciles reveals that the performance decline of automated underwriting systems is disproportionately concentrated among borrowers in the lowest income decile, while default rates for borrowers in the highest income decile show no change. This suggests that the deterioration in the relative performance of automated underwriting is most pronounced among borrowers severely affected by the pandemic, whose ability to repay loans was significantly diminished. Automated systems, reliant on historical data that was no longer representative of the new economic conditions, struggled to accurately predict defaults within this group. In contrast, human underwriters were able to adapt more readily to the rapidly changing environment by incorporating real-time information, resulting in comparatively better performance during this period.

To evaluate whether lenders adjusted their underwriting practices in response to the heightened economic uncertainty caused by the COVID-19 pandemic, I analyze changes in borrower attributes, loan contract terms, and the volume of loans originated, as any such adjustments could potentially explain the observed results. This analysis considers borrower attributes such as income, employment, or income verification, alongside key loan characteristics, including loan amount, maturity, interest rates, and loan-to-value ratio, during the post-pandemic period. Furthermore, I investigate whether the number of loans originated by automated systems shifts relative to those underwritten by humans after the onset of COVID-19. The results indicate no significant changes in these variables, suggesting that lenders did not modify their underwriting practices in response to the shock. This lack of adjustment may reflect the rigidity of automated underwriting systems, which depend on pre-established models and historical data, making rapid adaptation difficult. My findings are in line with Ben-David, J. Johnson, and Stulz 2025, who document that FinTech lenders continued to issue small business loans on essentially the same terms as before the pandemic, despite the dramatic shift in economic conditions. As they highlight, lenders were unable to assess how inaccurate their models had become under new conditions, since evaluating loan performance requires time to observe repayment behavior. Moreover, even if lenders attempted to revise their models in response to the shock, they lacked a credible way to test the reliability of such adjustments.

The validity of the difference-in-differences analysis relies on the assumption that, in the absence of the pandemic, the performance of loans originated by human underwriters and those by automated systems would have followed comparable trajectories over time. To assess this assumption, I take two steps. First, I limit the post-COVID analysis to three quarters to ensure that lenders had not yet made adjustments to their automation policies, either by updating their models or tightening their standards. Second, I examine pre-pandemic data on default rates for both underwriting methods. The results reveal no significant divergence in trends between the two groups prior to the pandemic, suggesting that both automated and human underwriting followed parallel paths in terms of default rates before the COVID-19 shock, which provides support for the validity of the parallel trends assumption.

My research contributes to several strands of literature. A significant body of work has examined the role of technology in finance, including the use of advanced machine learning models for credit scoring and default prediction (Fuster et al. 2022; Gambacorta et al. 2024; Khandani, Kim, and Lo 2010; Sadhwani, Giesecke, and Sirignano 2021), comparisons of realized delinquency rates between FinTech and traditional bank loans (Buchak et al. 2018; Fuster et al. 2019; Di Maggio, Ratnadiwakara, and Carmichael 2022), and analyses of whether humans or algorithms are more effective in screening and monitoring borrowers (Berg 2015;

Costello, Down, and Mehta 2020; Jansen, Nguyen, and Shams 2024; Gao, Yi, and Zhang 2024). However, these studies primarily focus on periods of economic stability. In this paper, I shift the focus to the performance of technology in loan origination during periods of economic instability, specifically the COVID-19 pandemic.

The study most closely related to mine is Jansen, Nguyen, and Shams 2024, which uses a randomized experiment to compare the performance of human and algorithmic underwriting in the U.S. auto loan market. Their findings show that algorithmic underwriting outperforms human underwriting in terms of higher loan profitability and lower default rates. However, their analysis is limited to stable economic conditions. My study complements their work by extending the analysis to a period of economic uncertainty—the COVID-19 pandemic. This extension is particularly important because it reveals the potential costs associated with the widespread automation of underwriting processes. While automation offers significant advantages during stable economic periods—such as increased efficiency, scalability, and profitability—my study highlights the risks and limitations of fully automating underwriting, particularly when unforeseen economic shocks occur. In the face of a crisis like the pandemic, automated models can struggle to adapt to rapidly changing conditions. As a result, the performance of automated systems in loan origination may deteriorate, leading to higher default rates and suboptimal loan decisions. This outcome underscores the crucial role of human expertise in times of economic uncertainty. Humans are better equipped to incorporate real-time information and make judgments based on current, non-quantifiable factors, which are often overlooked by automated models. My work emphasizes the trade-off between the benefits of automation and the risks of losing human expertise, suggesting that lenders may face significant costs if they rely solely on automated systems during crises.

Another closely related study is Ben-David, J. Johnson, and Stulz 2025, which examines the behavior of small-business FinTech lenders in the U.S. during March 2020, the onset of the COVID-19 crisis. Using detailed loan-level data from online lenders, they find that data-driven credit models used by these lenders fail to perform reliably when economic conditions deviate from those present in their training environment. They attribute the sharp contraction in credit supply by these lenders to model risk, as economic conditions deteriorated rapidly and the models became unreliable. This unreliability led lenders to reduce or halt lending. While their study focuses on supply-side decisions rather than the ex-post performance of these models, my paper complements their findings by examining loan outcomes, providing new evidence on how economic uncertainty affects the effectiveness of automated relative to manual underwriting.

The second strand of literature to which my paper contributes examines the impact of the COVID-19 shock on lending practices. Fuster et al. 2021 explore the role of FinTech lenders in the mortgage market during the pandemic, showing that automation facilitated the industry's ability to manage capacity constraints and operational challenges. Similarly, Bao and Huang 2021 compare FinTech and bank loans in China during the pandemic but focus on loans originated before COVID-19 and still active at the onset of the crisis. Their findings indicate that delinquency rates remained stable for bank loans but increased substantially for FinTech loans in the six months following the pandemic's start. My work differs from these studies in two key ways. First, I focus on within-lender comparisons of loans originated by automated systems versus those underwritten by humans, avoiding selection biases inherent in cross-firm comparisons of FinTech and non-FinTech institutions, which often serve distinct borrower profiles. Second, I analyze loans originated after the onset of the pandemic to examine how underwriting methods were directly affected by the crisis, rather than how pre-existing borrower characteristics shaped outcomes. This approach allows for a more precise evaluation of how the pandemic disrupted the predictive capabilities of automated systems compared to human decision-making, offering new insights into the adaptability of underwriting technologies during periods of economic uncertainty.

Third, my paper contributes to the growing literature on how technology shapes labor market outcomes and whether it replaces or augments human labor. While prior studies such as Brynjolfsson, Mitchell, and Rock 2018, Babina et al. 2023, Chen and Wang 2024, and

Kumar 2023 show that technology can have both augmenting and replacing effects depending on the task and context, my findings go further by identifying the conditions under which human labor remains essential. The underperformance of automated systems during periods of uncertainty highlights that the ability of technology to replace human labor depends not only on the nature of the task but also on the broader economic environment, particularly during times of disruption or crisis.

The remainder of the paper is structured as follows. Section 2 provides an overview of the institutional background. Section 3 describes the Reg AB II dataset. Section 4 details the empirical strategy employed in the analysis and presents the main results. Section 5 explores cross-sectional tests and potential mechanisms driving my results. Finally, Section 6 concludes.

2 Underwriting Process in Auto Loan Markets[1]

In indirect auto lending, which accounts for 90% of auto loans originated in the United States (Grunewald et al. 2023), the dealer acts as an intermediary between the borrower and the lender. The process begins when the customer selects a vehicle and negotiates the price, features, and options with the sales agent. Once the vehicle details are agreed upon, the customer then works with the Finance and Insurance (F&I) agent to arrange financing. Dealers typically have access to a broad network of lenders, allowing them to forward the customer's credit application to multiple financial institutions. The credit application, which includes essential personal and financial details such as the applicant's residential address, monthly income, mortgage or rent payments, Social Security Number (SSN), and other relevant information, is submitted electronically to the lenders along with the vehicle specifications and proposed loan terms.

Upon receiving a credit application, the lender generally obtains a credit report on the

1. This section partly draws on prospectuses issued by lenders in accordance with Regulation AB II.

applicant from one of the three national credit bureaus (Equifax, Experian, and TransUnion). The choice of bureau often depends on the lender's assessment of which bureau provides the most accurate and comprehensive credit report for the applicant's geographic area. If the applicant has sufficient credit history, the credit report will include the applicant's credit score, commonly referred to as the FICO score. Lenders also employ proprietary credit scoring algorithms developed by third-party credit scoring companies. These algorithms assign applicants a proprietary credit score, often referred to as a "custom credit risk score" or "scorecard." This score is used to assess the applicant's credit risk or creditworthiness based on the data provided by credit bureaus. Applicants are then categorized into tiers based on their credit risk and deal structure, which collectively determine their final pricing.

Applications are initially evaluated through an automated process. These applications are either automatically approved, automatically rejected, or forwarded for further review by a credit analyst. The automated process uses algorithms to assess applications based on various combinations of credit factors. Consequently, there are numerous clusters of credit factors that can lead to automated approval. These factors include FICO score, the lender's proprietary credit score, loan-to-value ratio, payment-to-income ratio, debt-to-income ratio, type of collateral, age of the collateral, and the mileage on the collateral, among others. Typically, applicants with a clean credit history, stable financial conditions, and a favorable deal structure are automatically approved. Conversely, applications characterized by higher risk—such as those with derogatory credit records, high debt burdens, low FICO scores, or high loan-to-value ratios—are automatically rejected.

A credit application is forwarded to a credit analyst in two scenarios: when credit-related terms fall outside the prescribed automatic approval thresholds, or when the application contains incomplete or inconsistent data, such as a mismatch in the SSN and address. Upon referral, the credit analyst is not provided with full visibility into the factors used by the automated system's algorithm to recommend a particular decision. Instead, the analyst evaluates the application based on the lender's established underwriting guidelines. The analyst considers key factors, including credit application data, credit bureau information, payment and debt ratios, and the applicant's prior experience with the lender, to reach a decision. In cases of incomplete or inconsistent data, the analyst may contact the dealer to verify and resolve the questionable information before proceeding. Based on their assessment of the strengths and weaknesses of each application, the credit analyst may approve or reject the application. Approval may be contingent upon specific conditions, such as the inclusion of a qualified co-applicant or guarantor, or adjustments to the loan terms, such as an increased down payment or a less expensive vehicle. In the final step, the underwriting decision is communicated to dealers electronically. The entire underwriting process in the U.S. auto loan markets is illustrated in Figure 1. Underwriting in the auto loan market differs from the mortgage market. While the heavily regulated mortgage market depends on standardized systems like DU and LP to meet GSE guidelines, the less regulated auto loan market allows lenders to use proprietary systems or third-party solutions.[2]

3 Data and sample

The data for this study are sourced from the Regulation AB II, created under the Dodd-Frank Act. This rule requires issuers of public auto loan asset-backed securities (ABS) to

2. Underwriting in the mortgage market is heavily influenced by strict regulatory requirements and the involvement of Government-Sponsored Enterprises (GSEs) such as Fannie Mae and Freddie Mac. These entities play a critical role in the secondary mortgage market by establishing standardized underwriting guidelines to ensure loan quality and consistency. Automated systems, including Fannie Mae's Desktop Underwriter (DU) and Freddie Mac's Loan Prospector (LP), evaluate loan applications against these guidelines. The outcomes of DU and LP classifications are either "Accept," indicating that the loan complies with GSE requirements and is eligible for purchase, or "Caution," which signals that further manual underwriting is necessary to determine eligibility (Johnson 2023).
Similarly, mortgage lenders seeking insurance from the Federal Housing Administration (FHA) must adhere to the FHA's underwriting guidelines. The U.S. Department of Housing and Urban Development (HUD) introduced the TOTAL (Technology Open to Approved Lenders) Mortgage Scorecard to standardize this process. TOTAL provides two possible classifications: "Accept" and "Refer." An "Accept" designation indicates that the borrower satisfies FHA underwriting requirements and is eligible for insurance, allowing the loan application to proceed. A "Refer" designation signifies that the system lacks sufficient information to make an automated determination, requiring a human underwriter to conduct a manual review and gather additional documentation to finalize the decision (Gao, Yi, and Zhang 2024).

report detailed loan data to the SEC every month, improving transparency in the ABS market (Sweet 2015). Momeni and Sovich 2022 compare the Regulation AB II dataset with the population of auto loans from Eqiafax Inc. and find that it closely reflects the characteristics of the broader U.S. auto loan market. This dataset provides comprehensive information, including loan, vehicle, and borrower characteristics at the time of origination, as well as performance histories throughout the life of each loan. A particularly valuable feature of this dataset is the inclusion of an underwriting indicator for each loan, enabling the identification of whether the loan was originated by an automated system or a human underwriter.

The U.S. auto loan market exhibits notable variation in the adoption of underwriting methods, with human judgment remaining a key component despite the advantages of automation. Figure 2 illustrates the differences in automation usage across lenders from January 2019 to September 2020. A select group of lenders relies entirely on automated underwriting systems, including Capital One, CarMax, Carvana, Ford Credit, Exeter Finance, Mercedes-Benz Financial, and World Omni. In contrast, several lenders employ a hybrid approach, integrating both automated and manual underwriting practices. This group includes AmeriCredit, BMW Financial Services, Honda Finance, Hyundai Motor Finance, Nissan Finance, Toyota Financial Services, USAA, and Volkswagen Financial Services. Notably, Fifth Third Bank stands out as the only lender in the sample to rely exclusively on manual underwriting, with no use of automation in its processes. These patterns underscore that human involvement in underwriting remains widespread, with manual processes continuing to play an important role in loan originations across the market. This reliance highlights the enduring value of human judgment in credit evaluation, even as automation becomes increasingly prevalent.

Figure 3 illustrates trends in the likelihood of automated underwriting over time. Panel A shows the average rate of automation across all lenders, demonstrating that it remains consistent and stable throughout the sample period. Panel B further breaks down the data into two categories: lenders with high automation rates (defined as those using automation for over 90% of their underwriting, represented by the blue dashed line) and lenders with lower automation rates (less than 90%, represented by the red solid line). Both groups display consistent trends, with no notable changes in automation practices observed even after the onset of the COVID-19 pandemic. These results indicate that the adoption and usage of automation remained steady despite the economic and operational challenges brought by the pandemic.

To test my hypothesis that lenders continue to rely on human underwriters because automated systems face limitations during periods of economic uncertainty, I restrict my estimation sample to auto loans originated one year before and three quarters after the onset of the COVID-19 pandemic (January 2019 to September 2020). This time frame enables an evaluation of the comparative performance of automated and manual underwriting systems under both stable and uncertain economic conditions. By focusing on the three quarters immediately following the onset of the pandemic, I ensure that lenders had not yet made adjustments to their underwriting practices, such as updating automated models or tightening credit standards. This restriction mitigates the risk of bias arising from policy changes that could otherwise confound the analysis of loan performance.

Further, I limit the sample to loans with complete data on critical variables, including interest rate, loan amount, loan maturity, scheduled monthly payment, vehicle condition (e.g., new or used), make-model-year, vehicle value, borrower credit score, and income. Additionally,I focus on lenders that employ both automated and manual underwriting methods, using within-lender variation to evaluate the performance of loans underwritten through different methods. I also remove loans with borrower income exceeding $250,000, vehicle values above $100,000, interest rates above 30 percent, and credit scores below 600.

Table 1 presents descriptive statistics for loans at the time of origination. The average loan in the sample has an interest rate of 5.8 percent, a scheduled monthly payment of $534, a vehicle price of $35,506, a maturity of 67 months, and an initial principal of $32,624. The average loan to value ratio is 90.3 percent. The borrowers in the sample have an average credit score of 756 and an annual household income of $102,100 The unconditional default rates are 1.3 percent at 24 months and 2.2 percent at 36 months.

The right-most columns of Table 1 provide a comparison between loans originated through automated systems (treated group) and those underwritten by human analysts (control group). For these comparisons, the sample is restricted to loans originated prior to the treatment date. Several notable differences emerge between the two groups. Loans originated through automated systems have higher average initial principals ($32,021 versus $28,928), longer maturities (67 months versus 65 months), higher interest rates (5 percent versus 3 percent), and smaller loan-to-value ratios (0.89 versus 0.94) compared to those underwritten by humans. Borrowers with automatically-underwritten loans also show higher average credit scores (759 versus 748) and greater household incomes ($104,000 versus $97,332) than borrowers with manually-underwritten loans.

To address the endogeneity in auto lenders' assignment of automated versus manual underwriting methods, I restrict my sample to lenders with identified discontinuities in the probability of automated underwriting. The quasi-random variation in the likelihood of automation allows for analyzing how the performance of automated versus human underwriting evolves before and after the onset of the COVID-19 pandemic. Using the data-driven procedures detailed in Section 4.2.1, I identify discontinuities in the probability of automated underwriting across various credit score thresholds for 2 captive lenders, Nissan Motor Acceptance Corporation and Volkswagen Financial Services. My final sample consists of 13,924 auto loans issued by these two lenders around the identified cutoffs. Table 2 shows descriptive statistics for these loans. The average loan in this restricted sample has an interest rate of 3.6 percent, a scheduled monthly payment of $513, a vehicle price of $31,133, a maturity of 68 months, and an initial principal of $30,305. The average loan to value ratio is 98.2 percent. The borrowers in this sample have an average credit score of 713 and an annual household income of $85,356 The unconditional default rates are 1.8 percent at 24 months and 2.9 percent at 36 months.

The right-most columns of Table 2 provide a comparison between loans originated through automated systems (treated group) and those underwritten by human analysts (control group) for the restricted sample. For these comparisons, the sample is restricted to loans originated prior to the treatment date. Several notable differences emerge between the two groups. Loans originated through automated systems have shorter maturities (66 months versus 68 months), lower interest rates (4 percent versus 4.3 percent), and smaller loan-to-value ratios (0.94 versus 0.99) compared to those underwritten by humans. Borrowers with automatically-underwritten loans also show higher average credit scores (714 versus 709) and lower household incomes ($84,826 versus $86,456) than borrowers with manually-underwritten loans.

Although automatically and manually-underwritten loans show observable time-invariant differences, the baseline difference-in-differences model in Section 4.1 addresses these differences by including fixed effects for lender, vehicle, and borrower characteristics. As shown in Figure 9, there is no evidence of differential pre-trends between the treated and control groups after accounting for these fixed effects. This indicates that, while the two groups differ in levels prior to treatment, their pre-treatment trajectories are identical.

4 Empirical Methodology

In the following section, I describe the empirical strategy used to analyze changes in the relative performance of loans originated through automated systems versus those underwritten by credit analysts in the period following the onset of the COVID-19 pandemic.

4.1 Difference-in-Differences (DiD) Design

To compare loan performance across underwriting methods before and after the onset of COVID-19, I employ a Difference-in-Differences (DiD) design. The treated group consists of loans originated through automated systems, while the control group includes loans underwritten manually by credit analysts. This approach compares the performance of these two groups in the periods before and after the onset of the pandemic. Formally, I estimate Equation 1:

$$y_{i,t} = \alpha + \gamma \cdot \text{Automated}_{i,t} + \beta \cdot \text{Automated}\times\text{After}_{i,t} + \delta_{s,t} + \delta_{l,t} + \delta_{c,t} + \delta_{I,t} + \delta_{v,t} + \eta_{i,t}, \quad (1)$$

where the outcome variable is 36-month default rate for loan i originated in quarter t. This specification includes fixed effects to control for lender ($\delta_{l,t}$) , state ($\delta_{s,t}$), vehicle ($\delta_{v,t}$), and borrowers' income ($\delta_{I,t}$) and credit score ($\delta_{c,t}$), and standard errors are clustered at the lender level. The coefficient of interest, β, captures the average change in default rates for automatically-underwritten loans relative to manually-underwritten loans after the onset of COVID-19.

Table 3 presents the coefficient estimates from estimating Equation 1. The coefficient of interest, β, is positive and statistically significant, indicating a notable increase in default rates for loans underwritten automatically compared to those underwritten manually in the post-COVID period. Furthermore, consistent with prior studies, γ is negative, indicating that, prior to the pandemic, loans originated by automated systems exhibited lower default rates, suggesting superior performance of automated underwriting compared to manual methods. However, these estimated coefficients from a traditional DiD framework are more likely to be biased since lenders have discretion over whether and to what extent they automate their underwriting processes.

4.2 Quasi-Random Variation in Automation Probability

To address this endogeneity in the assignment of automated and manual underwriting, I exploit quasi-random variation in the probability of automated underwriting, using lender-specific discontinuities in the likelihood of automation across various FICO score thresholds. By leveraging this variation, I narrow the comparison to loans originated just above and just below the identified thresholds. This allows me to isolate the performance differences between manually and automatically underwritten loans for borrowers with otherwise similar observable characteristics, thereby mitigating concerns about selection bias.

4.2.1 Detecting Discontinuities in Automation Probability

To detect discontinuities in the likelihood of automated underwriting, I estimate regressions where the dependent variable is an indicator for automated underwriting (1 for automated, 0 for manual) regressed on a set of indicator variables corresponding to 10-point FICO score bins for each lender. The 10-point bins start at FICO 600, with the first bin covering scores from 600 to 609, the second from 610 to 619, and so forth, up to scores of 850. The estimated coefficient for each FICO bin represents the average likelihood of automated underwriting for loans within each bin, relative to the omitted category (credit scores from 600 to 609).

This procedure identifies discontinuities in the likelihood of automation for two lenders: Nissan Motor Acceptance Corporation and Volkswagen Financial Services. As shown in Figures 4 and A.1, the results reveal that the probability of a loan being automatically underwritten increases as FICO scores rise. Specifically, Panel A of Figure 4 highlights a discontinuous jump in the unconditional likelihood of automation at a FICO score of 720 for Nissan, with a magnitude of 16 percentage points. Panel B of Figure 4 illustrates discontinuous jumps for Volkswagen at FICO score thresholds of 660 and 700, with magnitudes of 20 and 24 percentage points, respectively. As shown in Figure A.1, the jumps in the likelihood of automation are present even after conditioning on observable characteristics. This suggests that the locations of these discontinuities are randomly assigned and can generate quasi-random variation in the probability of automation.

4.2.2 First-Stage Regressions

To address the endogeneity concern, I use the locations of identified discontinuities as instruments for the two endogenous variables. Specifically, I formally estimate Equations 2 and 3:

$$\text{Automated}_{i,t} = \omega_1 + \Gamma_1 \cdot T_i + f(x_i) + T_i \cdot g(x_i) + \delta_{s,t} + \delta_{l,t} + \delta_{I,t} + \delta_{v,t} + \nu_{i,t}, \tag{2}$$

$$\text{Automated}\times\text{After}_{i,t} = \omega_2 + \Gamma_2 \cdot T_i \cdot \text{After}_t + f(x_i) + T_i \cdot g(x_i) \cdot \text{After}_t + \delta_{s,t} + \delta_{l,t} + \delta_{I,t} + \delta_{v,t} + \zeta_{i,t}, \tag{3}$$

where $\text{Automated}_{i,t}$ is a binary indicator equal to 1 if loan i is originated by an automated system and 0 otherwise. The instrument T_i is a dummy variable equal to 1 if the credit score is above the threshold and zero otherwise. The variable $\text{Automated}\times\text{After}_{i,t}$ captures the combined effect of automation and the post-pandemic period. The variable After_t is a dummy variable equal to 1 for loans originated in the three quarters following the onset of the COVID-19 pandemic and 0 otherwise. The forcing variable, x_i, is the centered credit score of loan i at origination, defined as the difference between the credit score and the lender specific threshold (credit score $-$ lender-specific threshold). Functions f and g are flexible functions of the centered credit score, capturing the relationship between the forcing variable and the outcome on either side of the threshold. The estimation sample is restricted to loans with credit scores within 20 points of an automated underwriting discontinuity.

4.2.3 Second-Stage Regressions

Next, I analyze how the COVID-19 pandemic affected the performance of loans originated by automated systems compared to those underwritten by credit analysts. To estimate the impact of the COVID-19 pandemic on loan default rates, I use the following baseline Difference-in-Differences (DiD) regression after instrumenting the endogenous variables:

$$y_{i,t} = \alpha + \gamma \cdot \widehat{\text{Automated}}_{i,t} + \beta \cdot \widehat{\text{Automated} \times \text{After}}_{i,t} + f(x_i) + T_i \cdot g(x_i) + \delta_{s,t} + \delta_{l,t} + \delta_{I,t} + \delta_{v,t} + \eta_{i,t}, \tag{4}$$

where the outcome variable is 36-month default rate for loan i originated in quarter t. $\widehat{\text{Automated}}_{i,t}$ and $\widehat{\text{Automated} \times \text{After}}_{i,t}$ are the predicted values from the first-stage regressions (Equations 2 and 3). f and g are flexible functions of the centered credit score to account for non-linear trends. This specification includes fixed effects to control for lender ($\delta_{l,t}$) , state ($\delta_{s,t}$), vehicle ($\delta_{v,t}$), and borrowers' income ($\delta_{I,t}$), and standard errors are clustered at the lender-FICO score level.

The coefficient of interest, β, captures the average change in default rates for automatically-underwritten loans relative to manually-underwritten loans after the onset of COVID-19. The analysis is conducted on a sample of auto loans originated by Nissan and Volkswagen between January 2019 and September 2020, focusing on loans just above and below the identified FICO score thresholds.

Table 4 presents the coefficient estimates from Equation 4. The coefficient of interest, β, is positive and statistically significant, indicating a notable increase in default rates for loans underwritten automatically compared to those underwritten manually in the post-COVID period. Furthermore, The coefficient on $\widehat{\text{Automated}}_{i,t}$,$\gamma$, captures the difference in default rates between automated and manually underwritten loans during the pre-pandemic period. This coefficient is negative, indicating that, prior to the pandemic, loans originated by automated systems exhibited lower default rates, suggesting superior performance of automated underwriting compared to manual methods. However, in the post-COVID period, this trend reverses, with the difference in default rates widening by approximately 630 basis points. The sum of these two coefficients yields a positive value, implying that, following the onset of the pandemic, the relative performance of loans underwritten by automated systems deteriorated compared to those underwritten manually. Specifically, the increase is 180 basis points, representing a 62 percent rise relative to the unconditional mean of 290 basis points.

These findings are robust to alternative model specifications, including varying fixed effects and polynomial orders. The results highlight the reduced adaptability of automated underwriting systems during periods of economic instability, as reflected in their diminished performance relative to manual underwriting methods in the wake of the COVID-19 pandemic.

To analyze how the impact of COVID-19 evolved over the sample period, I estimate Equation 5:

$$y_{i,t} = \alpha + \sum_{\tau=-4}^{2} \Gamma_\tau \cdot \widehat{\text{Automated}}_i \cdot \text{D}_\tau + \delta_a + \delta_{s,t} + \delta_{l,t} + \delta_{I,t} + \delta_{v,t} + \varepsilon_{i,t}, \tag{5}$$

where D_τ is equal to one if quarter t is τ quarters from the treatment date. I exclude the quarter prior to the treatment date, defined as the onset of COVID-19, ($\tau = -1$) as the reference quarter. Therefore, the Γ_τ coefficient captures the average difference in default rate between loans originated by automated systems and those underwritten manually in quarter τ relative to the average difference observed in default rate one quarter prior to the treatment date. I also control for the time-invariant variation in the likelihood of automation by adding δ_a. The results, shown in Figure 9, panels a and b, indicate a significant increase in default rates for automatically underwritten loans immediately after the onset of COVID-19. Furthermore, these figures provide strong evidence supporting the parallel trends assumption underlying the analysis. Specifically, there are no differential pre-trends in default rates between loans originated through automated systems and those underwritten by credit analysts prior to the pandemic.

4.2.4 Quasi-Random Variation: Identification Assumptions

In addition to the relevance condition (shown in Figures 4 and A.1), the internal validity of the fuzzy regression discontinuity (RD) design depends on the satisfaction of two identification assumptions. The first assumption is the exclusion restriction, which posits that crossing a credit score threshold influences the default rate only through its effect on the likelihood of automation and not through any other mechanism. The second assumption is local continuity, which requires that, in the absence of treatment, borrowers just below the FICO score threshold provide valid counterfactual for those just above it. Next, I provide empirical evidence to support these assumptions.

Exclusion Restriction: The exclusion restriction assumes that no other loan contract terms exhibit discontinuous changes around the credit score thresholds. To test this assumption, I estimate Equation 6:

$$y_{i,t} = \alpha + \gamma \cdot \widehat{\text{Automated}}_{i,t} + \beta \cdot \widehat{\text{Automated} \times} \text{After}_{i,t} + f(x_i) + T_i \cdot g(x_i) + \delta_{s,t} + \delta_{l,t} + \delta_{I,t} + \delta_{v,t} + \eta_{i,t}, \tag{6}$$

where $y_{i,t}$ represents the loan contract terms, including loan maturity, loan amount, loan to value ratio, and interest rate of loan i originated in quarter t . The forcing variable, x_i, is the centered credit score of loan i at origination, defined as the difference between the credit score and the lender specific threshold (credit score − lender-specific threshold). The instrument T_i is a dummy variable equal to 1 if the credit score is above the threshold. Functions f and g are flexible functions of the centered credit score, capturing the relationship between the forcing variable and the outcome on either side of the threshold. This specification includes fixed effects to control for lender ($\delta_{l,t}$) , state ($\delta_{s,t}$), vehicle ($\delta_{v,t}$), and borrowers' income ($\delta_{I,t}$).

Figures 5 and 6 graphically illustrate local continuity of loan terms across credit score thresholds for Nissan and Volkswagen respectively. These figures support the evidence that the locations of these discontinuities are randomly assigned. Next, I formally test whether the loan terms change around the thresholds in a regression analysis. Table 5 shows that estimated coefficients, γ, for loan maturity, loan amount, interest rate, and loan-to-value are statistically insignificant, providing no evidence of discontinuous changes in these variables at the thresholds. These findings support the validity of the exclusion restriction, as they suggest that the thresholds do not systematically affect these loan terms.

Next, I investigate whether the results are driven by any changes in the contractual terms of loans. In Table 5, estimated coefficients, β, show no significant changes in loan contract terms, providing robust evidence that variations in co-determined loan terms do not account for the observed deterioration in the performance of automated systems. This reinforces the conclusion that the primary findings are not confounded by shifts in loan contract characteristics. Furthermore, Figure 10 shows that loan contract terms do not systematically change after the onset of COVID-19 pandemic. These findings are consistent with Ben-David, J. Johnson, and Stulz 2025, who document that FinTech lenders continued to originate small business loans on essentially the same terms despite a sharp deterioration in economic conditions. They interpret this rigidity as evidence of model risk—that is, the inability of data-driven credit models to adjust quickly in response to rapidly evolving circumstances. Similarly, the stability of contract terms in my setting suggests that automated underwriting systems did not recalibrate in response to the COVID-19 shock, further supporting the interpretation that performance declines reflect limitations in risk assessment rather than shifts in contractual pricing or structure.

Local Continuity Assumption: The local continuity assumption implies that predetermined borrower characteristics should be similar on either side of the credit score threshold. To test this condition, I estimate Equation 6, using household income and the probability of having a borrower's income or employment being verified as the outcome variables. The results, presented in Table 6 indicate no change in these characteristics around the thresholds. The coefficient of interest, γ, is neither economically nor statistically distinguishable from zero, with the sole exception being Column 3. In this instance, the coefficient of $\text{Automated}_{i,t}$ achieves statistical significance at the 10% level. However, this result is not concerning, as the sign of the coefficient contradicts the expected finding. Figures 7 and 8 further illustrate the absence of discontinuities in these variables, providing robust support for the validity of the local continuity assumption.

Up to this point, I have focused on interpreting my results in terms of the intensive margin. However, shifts in borrower composition along the extensive margin may also have played a role in the decreased performance of automated loans during the post-COVID period. Although the inclusion of fixed effects helps to account for changes in borrower composition, it is crucial to determine whether the observed effects are primarily driven by the intensive margin or reflect broader shifts along the extensive margin. To analyze the impact of the pandemic on borrower composition, I change the focus of my analysis to the coefficient of $\text{Automated}\times\text{After}_{i,t}$, where β captures the average change in borrower characteristics for automatically underwritten loans relative to manually underwritten loans after the onset of the COVID-19 pandemic.

Table 6 present the coefficient of interest, γ. Consistent with the results indicating that the effects of COVID-19 operate primarily along the intensive margin, I find no significant deterioration in borrower characteristics for automated loans following the onset of the pandemic. The only exception is the borrower's income which achieves statistical significance at the 10% level. This is not concerning given that the sign of the coefficient contradicts the expected finding. Figure 11 illustrates the evolution of borrower characteristics around the treatment date. As before, the results show no significant declines in average household income or in the likelihood of a borrower's income or employment being verified. Furthermore, the absence of differential pre-trends across the treated and control groups supports the parallel trends assumption underlying the analysis.

Next, I investigate whether the onset of COVID-19 influenced the volume of loans originated by automated underwriting systems relative to manual underwriting. The results, presented in Table 7, yield two important insights. First, the volume of loans originated by automated systems is statistically indistinguishable from that of manually underwritten loans in pre-COVID period. This finding contrasts with prior evidence by Gao, Yi, and Zhang 2024, which suggest that automated systems outperform manual underwriting by originating a significantly greater number of loans. This discrepancy might be due to different samples. Gao, Yi, and Zhang 2024 show that automated underwriting leads to more credit access for high-risk borrowers during times of economic stability. My findings, however, show that for low-risk borrowers, the performance of automated and manual underwriting systems is not statistically different from each other. Second, the analysis reveals no significant change in the number of loans originated by either automated or manual underwriting systems following the onset of the pandemic. These findings indicate that variations in loan origination volumes, or the intensive margin, do not explain the observed results in this study. Taken together, these findings suggest that the composition of borrowers remained largely unchanged in response to COVID-19, reinforcing the conclusion that the observed deterioration in automated loan performance is not attributable to shifts in borrower characteristics.

5 Heterogenous Effects of COVID-19

In this section, I examine whether the observed increase in default rates is more pronounced among borrowers who were disproportionately affected by the economic disruptions of the COVID-19 pandemic. To explore this, I divide the sample into two groups based on the median (or decile) household income of borrowers. To capture the differential impacts of the pandemic across these income groups, I estimate the following triple-differences (DDD) model using Equation 7:

$$
\begin{aligned}
y_{i,t} = \alpha + \beta \cdot \text{Low income}_i \cdot \widehat{\text{Automated X}}\ \text{After}_{i,t} + \Gamma \cdot \widehat{\text{automatedXAfter}}_{i,t} \cdot \text{After}_t + \\
\theta \cdot \text{Low income}_i \cdot \widehat{\text{Automated}}_{i,t} + \delta_{s,t} + \delta_{l,t} + \delta_{I,t} + \delta_{v,t} + \varepsilon_{i,t}
\end{aligned}
\tag{7}
$$

where the outcome variable is 36-month default rate of loan i originated in quarter t. The variable Low credit score$_i$ is a dummy variable equal to 1 if the borrower's household income falls within the lowest income category and 0 if it falls within the highest income category. The coefficient of interest, β, measures the differential impact of the COVID-19 pandemic on the performance of loans originated by automated systems relative to those underwritten manually, specifically for borrowers in the lowest income category compared to borrowers in the highest income category.

If lower-income borrowers were more severely impacted by the pandemic, the historical data on which automated systems rely would fail to accurately predict their likelihood of default, leading to poorer performance for loans originated through automated systems. In contrast, credit analysts, who can directly account for pandemic-related disruptions such as job losses or reductions in income, are better equipped to evaluate borrowers in real time. This ability allows them to originate loans with improved performance outcomes, particularly for lower-income borrowers, compared to automated systems. Table 8 presents the coefficient estimates from the model. Consistent with expectations, the results indicate that the increase in default rates is significantly larger for loans originated by automated systems for borrowers in the lowest income category compared to those in the highest income category. Specifically, the performance of automated underwriting systems is statistically comparable to that of manual underwriters for very high-income borrowers. In contrast, their performance significantly deteriorates relative to human underwriters for very low-income borrowers. This disparity could be due to the fact that automated underwriting systems are less reliable or fail to accurately predict the default rates of individuals who are more likely to be adversely affected by this unexpected large scale shock.

6 Conclusion

In this study, I investigate the performance of automated versus human underwriting in the U.S. auto loan market during the COVID-19 pandemic, addressing the question of why some lenders continue to rely on human underwriters despite the widely recognized advantages of automation. While automated systems are acknowledged for their efficiency, profitability, and scalability during stable economic conditions, the varying levels of adoption across lenders suggests a trade-off between these benefits and the adaptability offered by human decision-making. My findings provide critical insights into this trade-off, particularly in the context of economic uncertainty.

Using the COVID-19 pandemic as a natural experiment, I examine how unexpected shocks impact the relative effectiveness of automated and human underwriting. My findings support the hypothesis that human underwriters are better able to adapt to unprecedented conditions. Specifically, automated systems, reliant on historical data and pre-established models, struggled to respond to the sudden and unpredictable economic changes caused by the pandemic. This rigidity led to higher default rates for loans originated by automated systems compared to those underwritten by humans. In contrast, human underwriters, leveraging real-time information and contextual judgment, demonstrated superior adaptability, particularly in assessing borrowers most vulnerable to the economic disruptions, such as those in lower-income groups.

The results also reveal that lenders relying solely on automated systems face significant risks during economic shocks, as these systems are unable to recalibrate quickly enough to account for rapidly changing borrower conditions. This limitation underscores the importance of maintaining human expertise as part of the underwriting process, particularly during periods of economic uncertainty. While automation delivers substantial benefits in stable environments, its constraints during crises highlight the need for a balanced approach that integrates the efficiency of technology with the adaptability of human decision-making.

References

Babina, Tania, Anastassia Fedyk, Alex X He, and James Hodson. 2023. *Firm investments in artificial intelligence technologies and changes in workforce composition.* Vol. 31325. National Bureau of Economic Research.

Bao, Zhengyang, and Difang Huang. 2021. “Shadow banking in a crisis: Evidence from FinTech during COVID-19.” *Journal of Financial and Quantitative Analysis* 56 (7): 2320–2355.

Ben-David, Itzhak, Mark J. Johnson, and René M Stulz. 2025. “Models behaving badly: The limits of data-driven lending.” *Review of Finance,* rfaf009.

Berg, Tobias. 2015. “Playing the devil’s advocate: The causal effect of risk management on loan quality.” *The Review of Financial Studies* 28 (12): 3367–3406.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018. “What can machines learn and what does it mean for occupations and the economy?” In *AEA papers and proceedings,* 108:43–47. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Buchak, Greg, Gregor Matvos, Tomasz Piskorski, and Amit Seru. 2018. “Fintech, regulatory arbitrage, and the rise of shadow banks.” *Journal of financial economics* 130 (3): 453–483.

Chen, Mark A, and Joanna Wang. 2024. “Displacement or Augmentation? The Effects of AI Innovation on Workforce Dynamics and Firm Value.”

Costello, Anna M, Andrea K Down, and Mihir N Mehta. 2020. “Machine+ man: A field experiment on the role of discretion in augmenting AI-based lending models.” *Journal of Accounting and Economics* 70 (2-3): 101360.

Di Maggio, Marco, Dimuthu Ratnadiwakara, and Don Carmichael. 2022. *Invisible primes: Fintech lending with alternative data.* Technical report. National Bureau of Economic Research.

Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther. 2022. "Predictably unequal? The effects of machine learning on credit markets." *The Journal of Finance* 77 (1): 5–47.

Fuster, Andreas, Aurel Hizmo, Lauren Lambie-Hanson, James Vickery, and Paul S Willen. 2021. *How resilient is mortgage credit supply? Evidence from the COVID-19 pandemic.* Technical report. National Bureau of Economic Research.

Fuster, Andreas, Matthew Plosser, Philipp Schnabl, and James Vickery. 2019. "The role of technology in mortgage lending." *The Review of Financial Studies* 32 (5): 1854–1899.

Gambacorta, Leonardo, Yiping Huang, Han Qiu, and Jingyi Wang. 2024. "How do machine learning and non-traditional data affect credit scoring? New evidence from a Chinese fintech firm." *Journal of Financial Stability,* 101284.

Gao, Janet, Hanyi Livia Yi, and David Zhang. 2024. "Algorithmic Underwriting in High Risk Mortgage Markets."

Grunewald, Andreas, Jonathan A. Lanning, David C. Low, and Tobias Salz. 2023. "Auto Dealer Loan Intermediation: Consumer Behavior and Competitive Effects." Available on authors' websites.

Howell, Sabrina T, Theresa Kuchler, David Snitkof, Johannes Stroebel, and Jun Wong. 2024. "Lender automation and racial disparities in credit access." *The Journal of Finance* 79 (2): 1457–1512.

Jansen, Mark, Hieu Quang Nguyen, and Amin Shams. 2024. "Rise of the machines: The impact of automated underwriting." *Management Science.*

Johnson, Stephanie. 2023. "Financial technology and the 1990s housing boom." *Available at SSRN 3619105.*

Khandani, Amir E, Adlar J Kim, and Andrew W Lo. 2010. "Consumer credit-risk models via machine-learning algorithms." *Journal of Banking & Finance* 34 (11): 2767–2787.

Kumar, Ishitha. 2023. "Impact of Robo-advisors on the Labor Market for Financial Advisors." *Available at SSRN 4608403.*

Momeni, Morteza, and David Sovich. 2022. "The marginal propensity to default." Working paper.

Sadhwani, Apaar, Kay Giesecke, and Justin Sirignano. 2021. "Deep learning for mortgage risk." *Journal of Financial Econometrics* 19 (2): 313–368.

Sweet, Charles A. 2015. "The SEC Finally Adopts Regulation AB II." *Journal of Structured Finance* 20 (4): 22.

Table 1: Descriptive statistics

	(1) Mean	(2) SD	(3) P10	(4) P25	(5) P50	(6) P75	(7) P90	(8) Automated	(9) Manual
Automated	0.746	0.435	0	0	1	1	1		
Subvention	0.920	0.271	1	1	1	1	1	0.908	0.891
Co-Obligor	0.310	0.463	0	0	0	1	1	0.284	0.374
Monthly payment	534	208	314	394	498	635	795	545	508
Credit score	756	65	659	712	764	811	835	759	748
Loan amount	31,624	12,816	18,150	22,998	29,232	37,776	48,161	32,021	28,928
Maturity	67	9	60	61	72	73	74	67	65
Car value	35,506	12,533	22,799	26,521	31,964	42,136	52,413	36,648	31,167
Income	102,100	69,817	40,008	55,221	82,800	124,716	186,000	104,000	97,332
Loan-to-Value	0.903	0.216	0.604	0.772	0.929	1.060	1.167	0.885	0.939
Interest rate	0.058	0.300	0	0.019	0.030	0.050	0.074	0.080	0.063
12-month default	0.004	0.065	0	0	0	0	0	0.005	0.003
24-month default	0.013	0.113	0	0	0	0	0	0.014	0.011
36-month default	0.022	0.146	0	0	0	0	0	0.022	0.019

NOTE.—This table describes my sample of 312,410 auto loans. Descriptive statistics are as of the loan origination date. In Columns 8 and 9, I compare auto loans that were automatically underwritten during the pre-treatment period to loans that were manually underwritten.

Table 2: Descriptive statistics (around the cutoffs)

	(1) Mean	(2) SD	(3) P10	(4) P25	(5) P50	(6) P75	(7) P90	(8) Automated	(9) Manual
Automated	0.323	0.468	0	0	0	1	1		
Subvention	0.915	0.279	1	1	1	1	1	0.894	0.901
Co-Obligor	0.243	0.429	0	0	0	0	1	0.173	0.28
Monthly payment	513	224	306	377	463	582	780	526	529
Credit score	713	20	688	705	716	728	735	714	709
Loan amount	30,305	11,965	18,798	22916	27,902	34,401	44,834	30,208	30,563
Maturity	68	7	60	61	72	72	73	66	68
Car value	31,133	11,096	20,890	24,105	28,099	34,522	44,485	32,492	31,132
Income	85,356	62,506	34,995	47,617	68,000	101,000	151,000	84,826	86,456
Loan-to-Value	0.982	0.198	0.721	0.865	1.005	1.132	1.224	0.939	0.989
Interest rate	0.036	0.027	0	0.009	0.039	0.057	0.072	0.040	0.043
12-month default	0.006	0.076	0	0	0	0	0	0.010	0.024
24-month default	0.018	0.132	0	0	0	0	0	0.029	0.037
36-month default	0.029	0.169	0	0	0	0	0	0.04	0.053

NOTE.—This table describes my sample of 13,924 auto loans around the identified cutoffs. Descriptive statistics are as of the loan origination date. In Columns 8 and 9, I compare auto loans that were automatically underwritten during the pre-treatment period to loans that were manually underwritten.

Table 3: Difference-in-differences regression (all loans): loan performance

	(1)	(2)	(3)	(4)	(5)
	36-month default	36-month default	36-month default	36-month default	36-month default
Automated	-0.014***	-0.012***	-0.006**	-0.005**	-0.012***
	(-3.72)	(-3.67)	(-2.22)	(-2.15)	(-3.44)
Automated X After	0.004	0.004	0.004**	0.004**	0.005*
	(1.46)	(1.45)	(2.40)	(2.55)	(1.93)
Control					Yes
LenderXTime FE	Yes	Yes	Yes	Yes	Yes
VehicleXTime FE		Yes	Yes	Yes	Yes
StateXTime FE			Yes	Yes	Yes
IncomeXTime FE			Yes	Yes	Yes
Credit ScoreXTime FE			Yes	Yes	Yes
MaturityXTime FE				Yes	Yes
R^2	0.092	0.120	0.131	0.131	0.114
Obs	312,410	312,341	312,340	312,340	259,032

NOTE.—This table reports coefficient estimates from Equation 1. The dependent variable is an indicator variable for whether the borrower was in default 36 months after originating their loan. The sample is restricted to auto loans originated a year before the treatment date and 3 quarters after the treatment date. Standard errors, presented below the coefficient estimates, are clustered at the lender level.

* Significant at the 10% level.

** Significant at the 5% level.

*** Significant at the 1% level.

Table 4: Difference-in-differences regression: loan performance

	(1) 36-month default	(2) 36-month default	(3) 36-month default	(4) 36-month default	(5) 36-month default
Automated	-0.072	-0.063	-0.053	-0.045	-0.045
	(-1.26)	(-1.05)	(-0.83)	(-0.74)	(-0.74)
Automated X After	0.081***	0.073***	0.069***	0.062**	0.063**
	(3.21)	(2.63)	(2.60)	(2.43)	(2.46)
Running	-0.001	-0.001	-0.001	-0.001	-0.001
	(-0.79)	(-0.79)	(-0.91)	(-0.82)	(-0.82)
Running X Treat	0.000	0.000	0.001	0.000	0.000
	(0.58)	(0.31)	(0.97)	(0.76)	(0.74)
$Running^2$ X Treat	0.000	0.000	0.000	0.000	0.000
	(0.63)	(0.72)	(0.48)	(0.56)	(0.57)
Control					Yes
LenderXTime FE	Yes	Yes	Yes	Yes	Yes
VehicleXTime FE		Yes	Yes	Yes	Yes
StateXTime FE			Yes	Yes	Yes
IncomeXTime FE			Yes	Yes	Yes
MaturityXTime FE				Yes	Yes
Obs	13,924	13,893	13,864	13,859	13,859

NOTE.—This table reports coefficient estimates from Equation 4. The dependent variable is an indicator variable for whether the borrower was in default 36 months after originating their loan. The sample is restricted to auto loans originated a year before the treatment date and 3 quarters after the treatment date.Standard errors, presented below the coefficient estimates, are clustered at the lender-credit score bin level.

* Significant at the 10% level.

** Significant at the 5% level.

*** Significant at the 1% level.

Table 5: Difference-in-differences regression: loan terms

	(1) Loan-to-Value	(2) log(maturity)	(3) log(loan amount)	(4) Interest rate
Automated	0.024	0.003	0.050	-0.025
	(0.30)	(0.81)	(0.58)	(-1.31)
Automated X After	0.027	-0.003**	0.000	0.011
	(0.53)	(-2.51)	(0.00)	(0.97)
Running	-0.001***	-0.000	-0.001***	-0.000***
	(-4.63)	(-1.30)	(-3.09)	(-8.76)
Running X Treat	-0.000	-0.000	-0.001	0.000
	(-0.16)	(-0.44)	(-0.87)	(1.24)
$Running^2$ X Treat	0.000	0.000	0.000	-0.000
	(0.99)	(0.74)	(0.90)	(-1.11)
Control	Yes	Yes	Yes	Yes
LenderXTime FE	Yes	Yes	Yes	Yes
VehicleXTime FE	Yes	Yes	Yes	Yes
StateXTime FE	Yes	Yes	Yes	Yes
IncomeXTime FE	Yes	Yes	Yes	Yes
MaturityXTime FE	Yes	Yes	Yes	Yes
Obs	13,859	13,859	13,859	13,859

NOTE.—This table reports coefficient estimates from Equations 4. The dependent variable is either the loan to value ratio, the natural log of the maturity, the natural log of the loan amount, or the interest rate. Standard errors, presented below the coefficient estimates, are clustered at the lender-credit score level.

* Significant at the 10% level.

** Significant at the 5% level.

*** Significant at the 1% level.

Table 6: Difference-in-differences regression: borrower composition

	(1) Employment verification	(2) Income verification	(3) Income
Automated	-0.015	-0.015	-0.043*
	(-0.51)	(-0.49)	(-1.72)
Automated X After	0.016	0.015	0.036*
	(0.60)	(0.58)	(1.68)
Running	-0.000	-0.000	0.000
	(-1.18)	(-1.18)	(0.28)
Running X Treat	0.000	0.000	0.001
	(0.68)	(0.68)	(1.08)
Running2 X Treat	0.000	0.000	-0.000
	(0.58)	(0.60)	(-1.12)
Control	Yes	Yes	Yes
LenderXTime FE	Yes	Yes	Yes
VehicleXTime FE	Yes	Yes	Yes
StateXTime FE	Yes	Yes	Yes
IncomeXTime FE	Yes	Yes	Yes
MaturityXTime FE	Yes	Yes	Yes
R^2	-0.001	-0.001	-0.019
Obs	13,859	13,859	13,859

NOTE.—This table reports coefficient estimates from Equations 4. The dependent variable is either an indicator variable for whether the borrower's employment is verified, the borrower's income is verified, or the natural log of the borrower's income. Standard errors, presented below the coefficient estimates, are clustered at the lender-credit score level.

* Significant at the 10% level.

** Significant at the 5% level.

*** Significant at the 1% level.

Table 7: Difference-in-differences regression: loan originations

	(1) Log(1+Number)	(2) Log(1+Number)
Automated	0.005	0.006
	(0.82)	(0.77)
AutomatedXAfter	-0.009	-0.010
	(-0.97)	(-1.13)
Lender FE	Yes	Yes
StateXTime FE	Yes	Yes
IncomeXTime FE		Yes
Credit ScoreXTime FE		Yes
Obs	93,219	93,219

NOTE.—This table reports coefficient estimates from Equation 4. The dependent variable is the log of one plus the number of loans originated. I estimate a linear regression model for these specifications. I calculate the number of loan originations at the lender × state × vehicle × income × origination year- quarter level. Standard errors, presented below the coefficient estimates, are clustered at the lender-credit score bin level.

* Significant at the 10% level.

** Significant at the 5% level.

*** Significant at the 1% level.

Table 8: Triple-differences regression: Income

	(1) 36-month default	(2) 36-month default
Automated X After	0.050	-0.118
	(1.51)	(-1.03)
AutomatedXAfterXLowIncome (median)	0.041	
	(1.22)	
AutomatedXAfterXLowIncome (decile)		0.280*
		(1.74)
Running	-0.001**	-0.001*
	(-2.35)	(-1.85)
Running X Treat	0.000	0.003
	(0.14)	(0.73)
Running$^2 XTreat$	0.000	-0.000
	(0.63)	(-0.38)
LenderXTime FE	Yes	Yes
VehicleXTime FE	Yes	Yes
StateXTime FE	Yes	Yes
Obs	13,792	2,534

NOTE.—This table reports coefficient estimates from Equation 7. The dependent variable is an indicator variable for whether the borrower was in default 36 months after originating their loan. In column (1), the indicator variable Low income is equal to one if the borrower's income is in the first quartile and zero if the borrower's income is in the fourth quartile. In column (2), the indicator variable Low income is equal to one if the borrower's income is in the first decile and zero if the borrower's income is in the tenth decile. Standard errors, presented below the coefficient estimates, are clustered at the lender-credit score bin level.

* Significant at the 10% level.

** Significant at the 5% level.

*** Significant at the 1% level.

Figure 1: Underwriting process in auto loan market

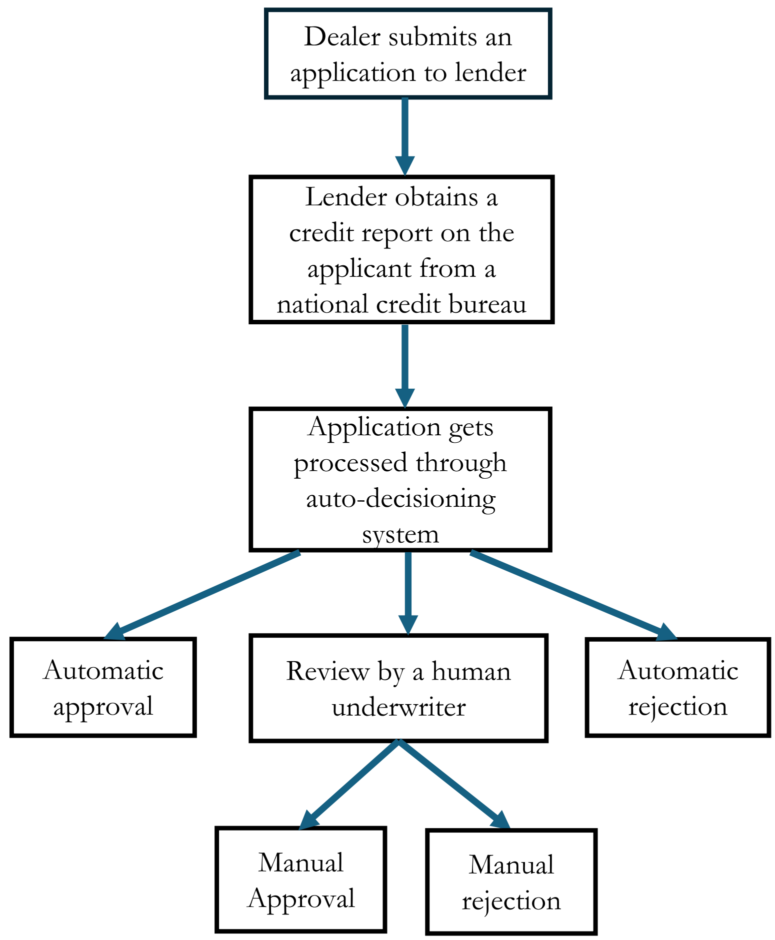

NOTE.——This figure presents automated underwriting process in the U.S. auto loan market.

Figure 2: Distribution of loans across lenders

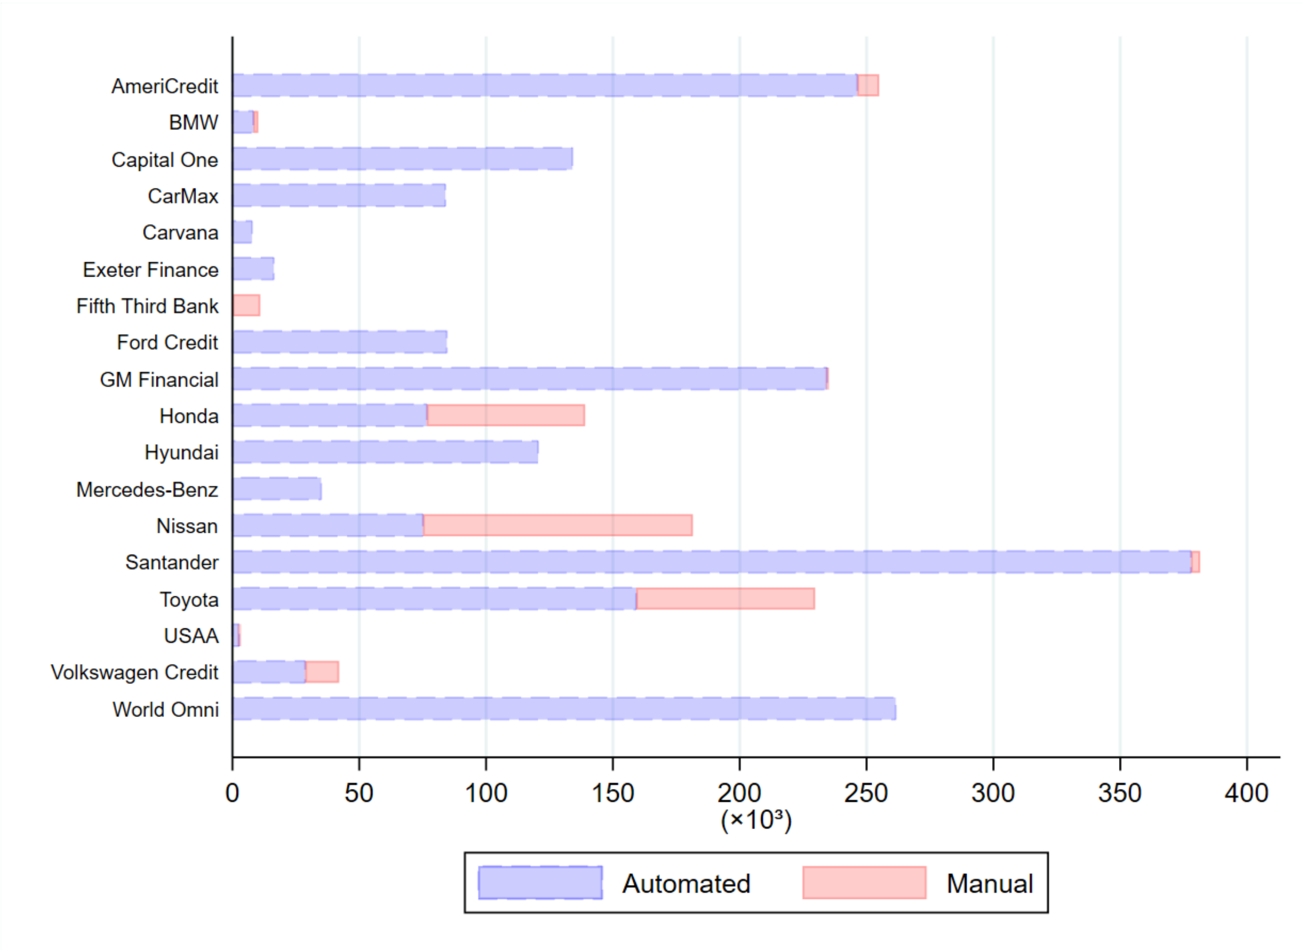

NOTE.——This figure plots the distribution of loans across lenders from January 2019 to September 2020. The blue box shows the number of loans that are automatically originated. The red box shows the number of loans that are manually originated

Figure 3: Heterogeneity in automated underwriting

(a)

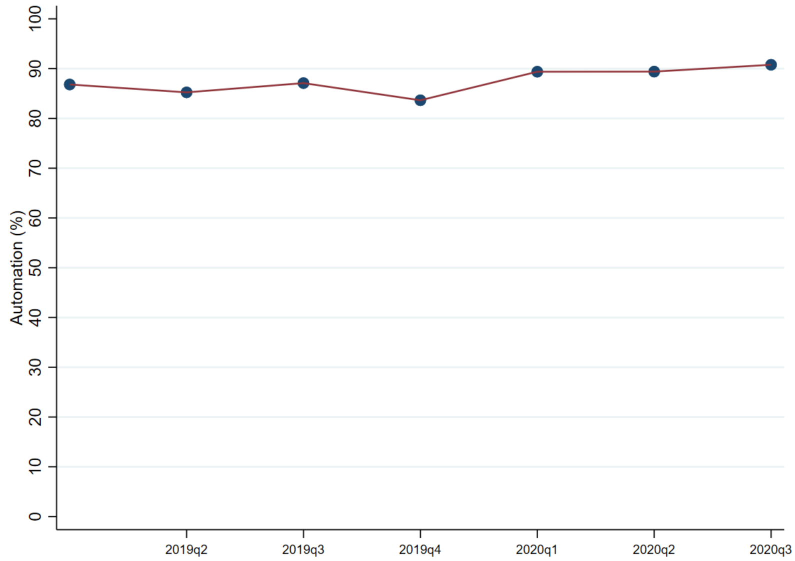

(b)

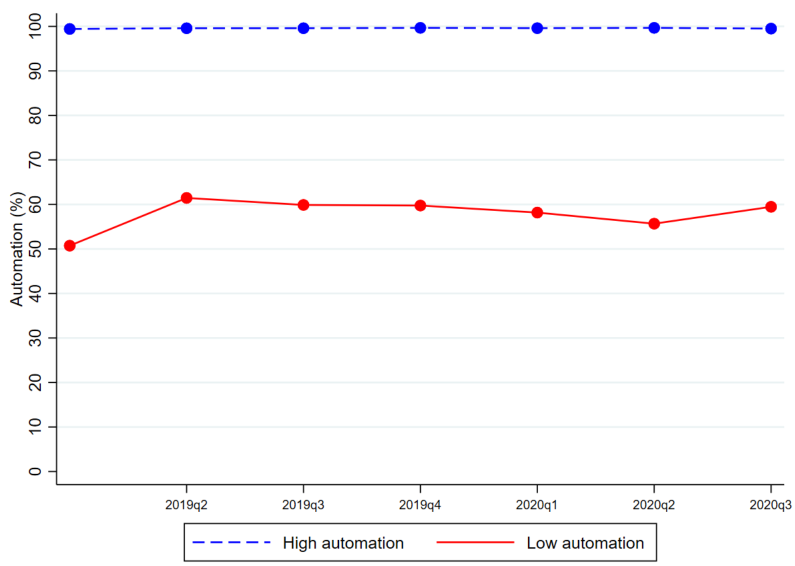

NOTE.——These figures plot the heterogeneity in automated underwriting over time. Panel (a) shows the average likelihood of automated underwriting over time. Panel (b) shows the average likelihood of automated underwriting over time for lenders with a high rate of automation (>90%) or a low rate of automation (< 90%).

Figure 4: Discontinuities in probability of automated underwriting

(a) Nissan

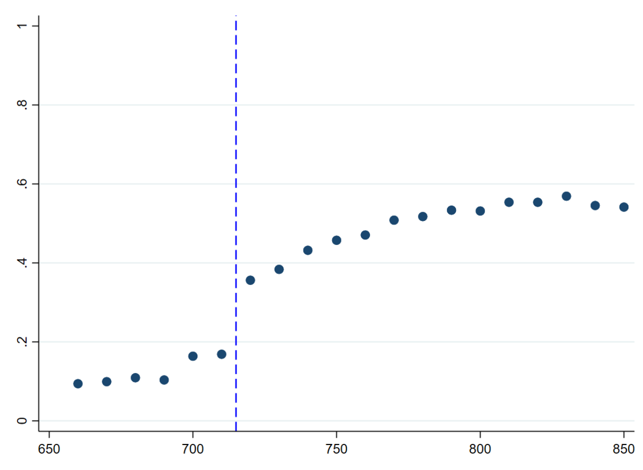

(b) Volkswagen

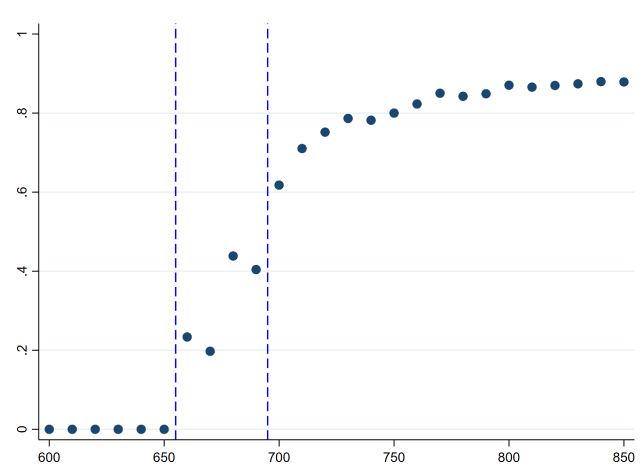

NOTE.——This figure plots the probaility of automated underwriting across credit scores for (a) Nissan and (b) Volkswagen.

Figure 5: Local continuity of loan terms: Nissan

(a) Loan-to-value

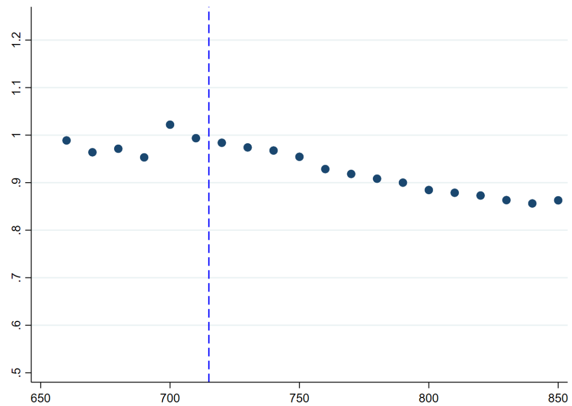

(b) Loan maturity

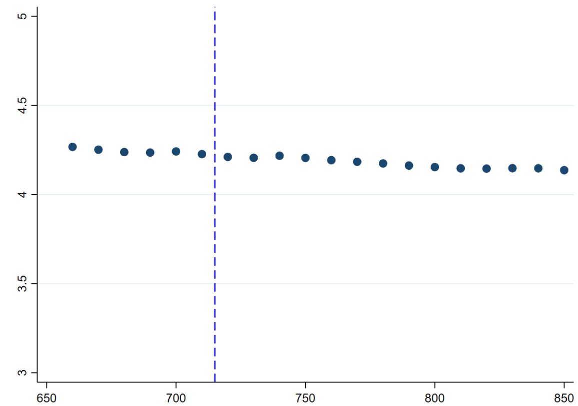

(c) Loan amount

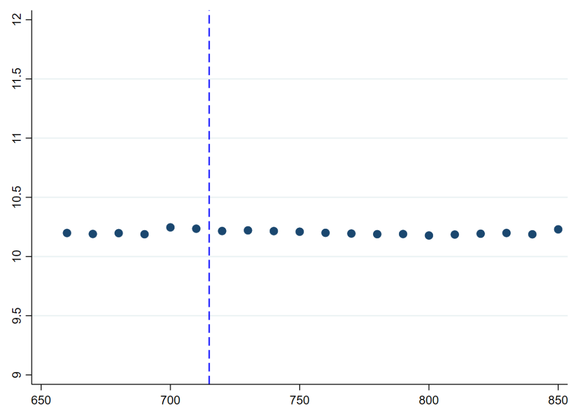

(d) Interest rate

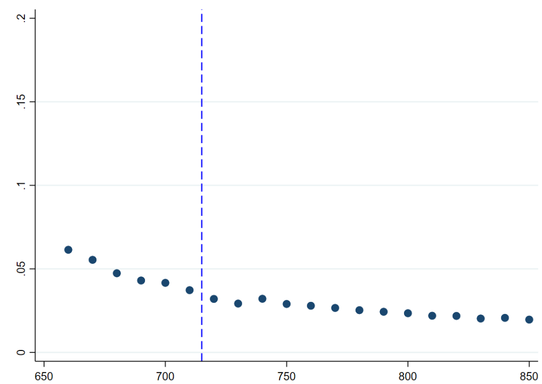

NOTE.——This figure plots raw average of loan terms across credit scores for Nissan. The loan term is either (a) loan-to-value ratio, (b) the natural log of the loan maturity, (c) the natural log of the loan amount, or (d) interest rate. The circles correspond to the average value of each loan term at a 10-point credit score bin.

Figure 6: Local continuity of loan terms: Volkswagen

(a) Loan-to-value

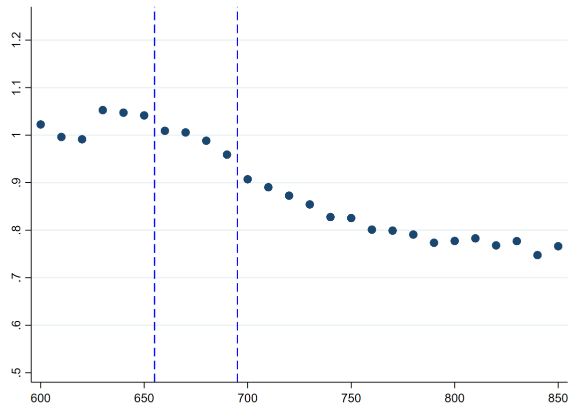

(b) Loan maturity

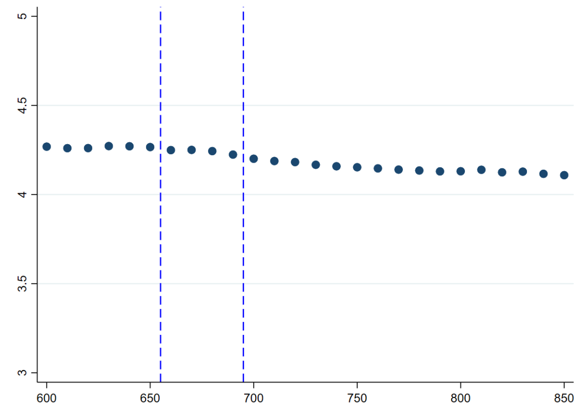

(c) Loan amount

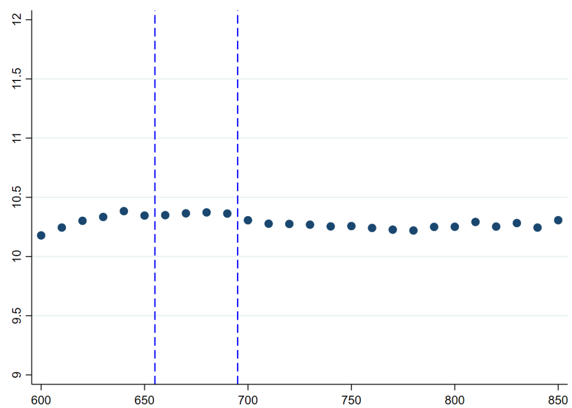

(d) Interest rate

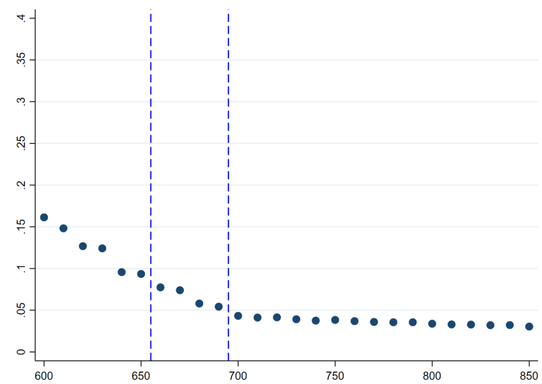

NOTE.——This figure plots raw average of loan terms across credit scores for Volkswagen. The loan term is either (a) loan-to-value ratio, (b) the natural log of the loan maturity, (c) the natural log of the loan amount, or (d) interest rate. The circles correspond to the average value of each loan term at a 10-point credit score bin.

Figure 7: Local continuity of borrower characteristics: Nissan

(a) Income

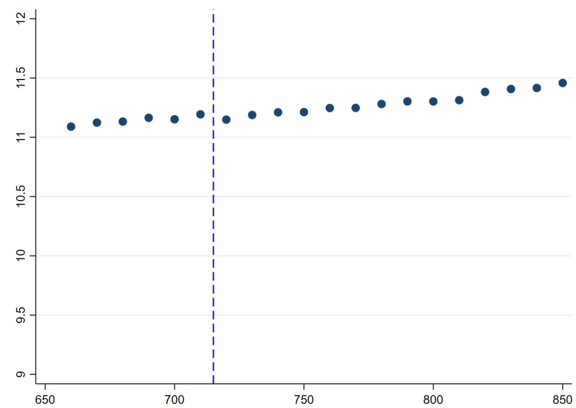

(b) Income verification

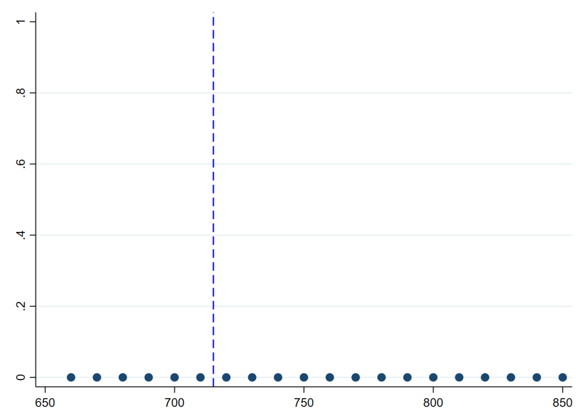

(c) Employment verification

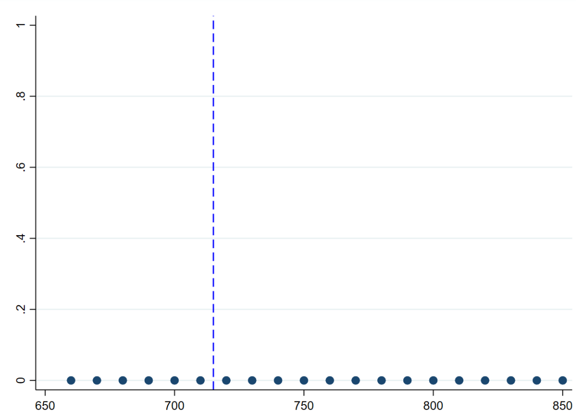

NOTE.——This figure plots raw average of borrower characteristics across credit scores for Nissan. The characteristic is either (a) the natural log of the income, (b) a borrower's income is verified, or (c) a borrower's employment is verified. The circles correspond to the average value of each borrower characteristic at a 10-point credit score bin.

Figure 8: Local continuity of borrower characteristics: Volkswagen

(a) Income

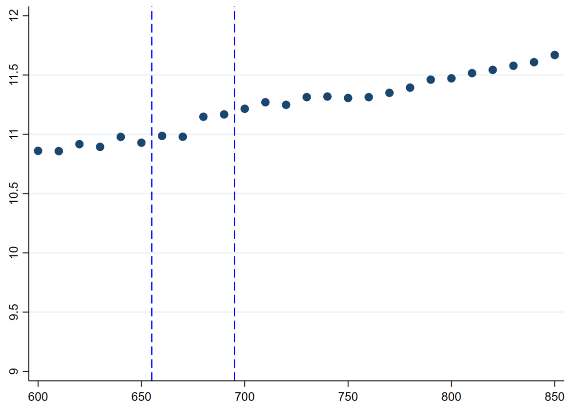

(b) Income verification

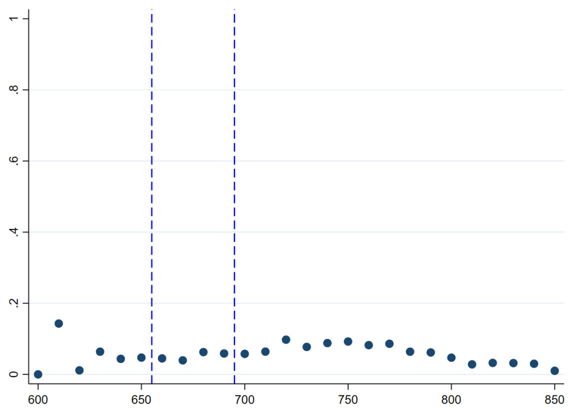

(c) Employment verification

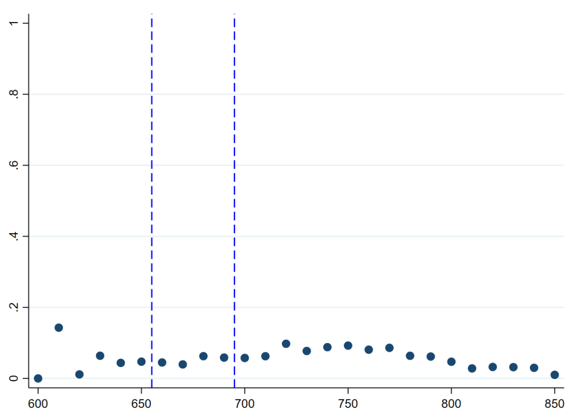

NOTE.——This figure plots raw average of borrower characteristics across credit scores for Volkswagen. The characteristic is either (a) the natural log of the income, (b) a borrower's income is verified, or (c) a borrower's employment is verified. The circles correspond to the average value of each borrower characteristic at a 10-point credit score bin.

Figure 9: Baseline specification: loan performance

(a) 24-month default

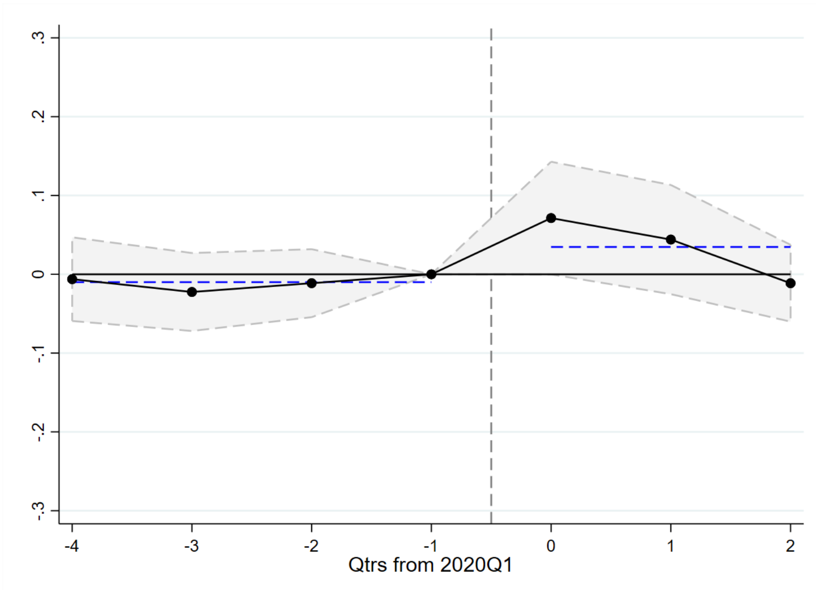

(b) 36-month default

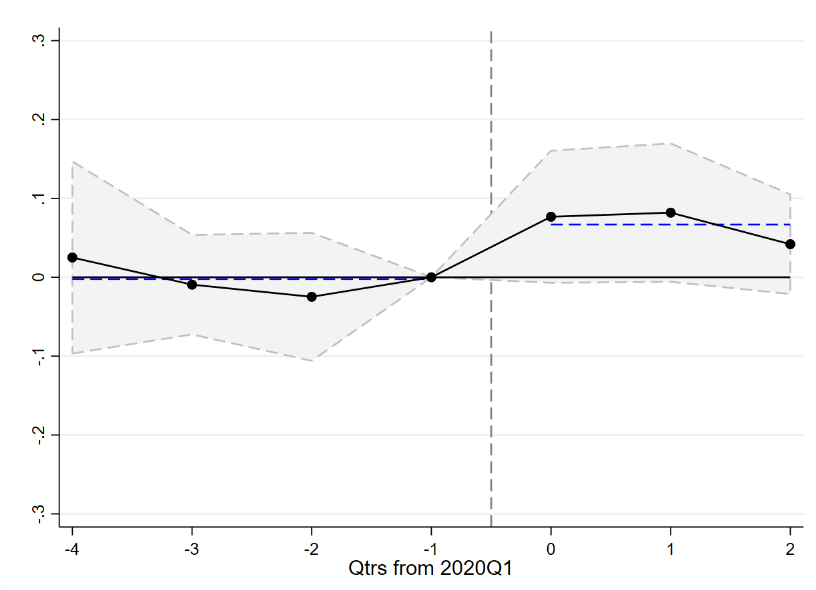

NOTE.——This figure plots coefficient estimates from Equation 5. The dependent variable is either the 24-month default or 36-month default. The x-axis corresponds to the number of quarters from the treatment date. The quarter $\tau = -1$ is the reference quarter. The circles correspond to the coefficient estimates, and the vertical bars correspond to 95 percent confidence intervals. Standard errors are clustered at the lender-credit score bin level.

Figure 10: Baseline specification: loan terms

(a) Loan-to-Value

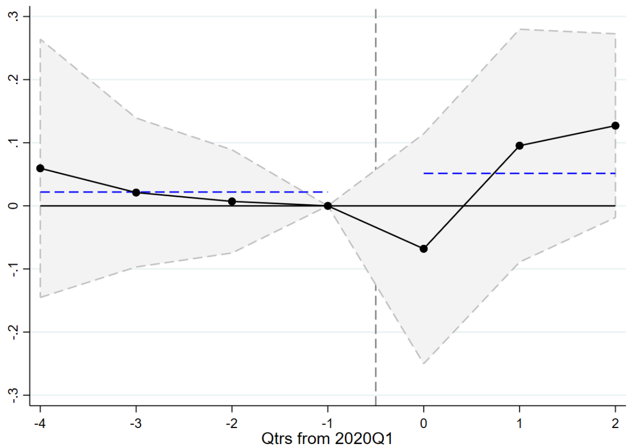

(b) Maturity

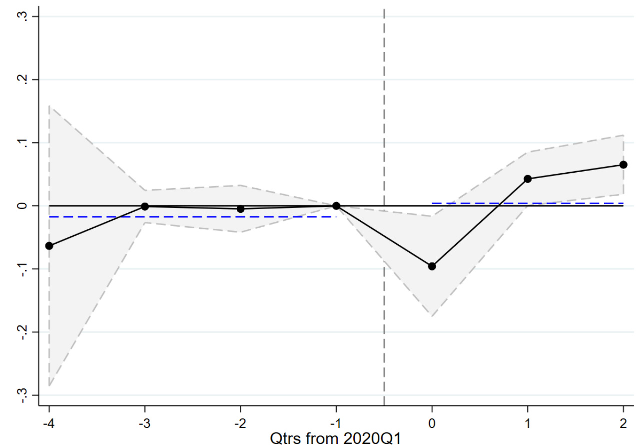

(c) Loan amount

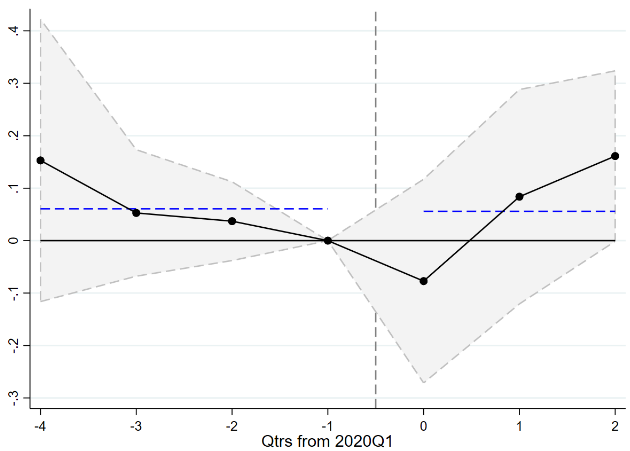

(d) Car value

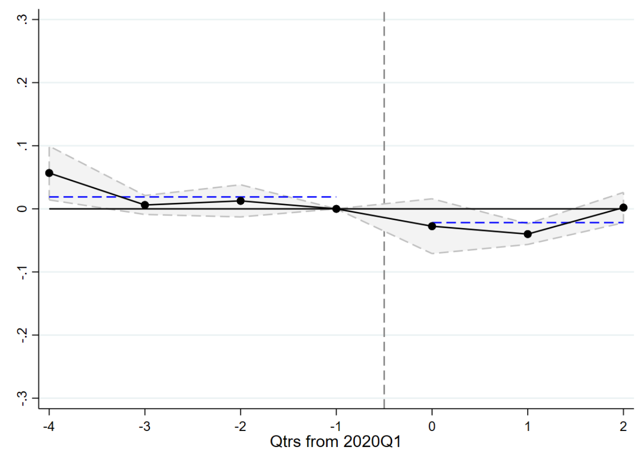

NOTE.——This figure plots coefficient estimates from Equation 5. The loan term is either (a) loan-to-value ratio, (b) the natural log of the loan maturity, (c) the natural log of the loan amount, or (d) the natural log of the car value. The x-axis corresponds to the number of quarters from the treatment date. The quarter $\tau = -1$ is the reference quarter. The circles correspond to the coefficient estimates, and the vertical bars correspond to 95 percent confidence intervals. Standard errors are clustered at the lender-credit score bin level.

Figure 11: Baseline specification: borrower characteristics

(a) Income

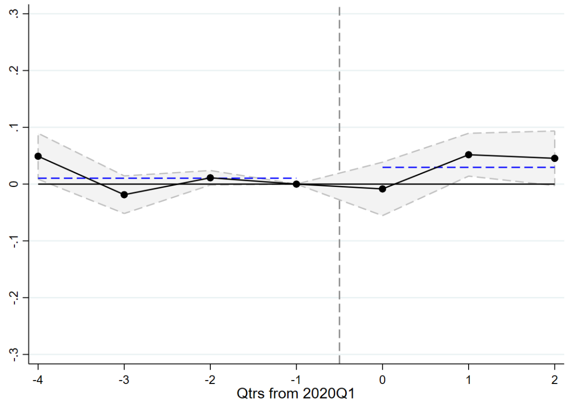

(b) Income verification

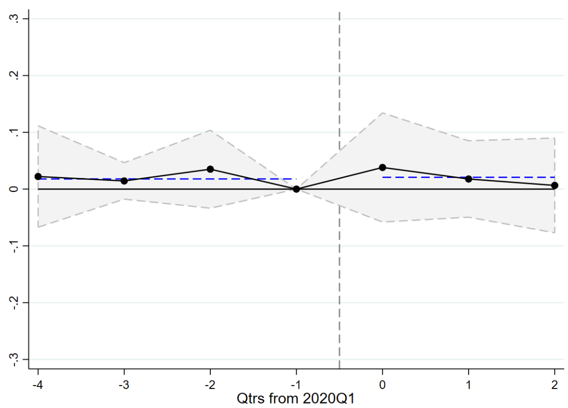

(c) Employment verification

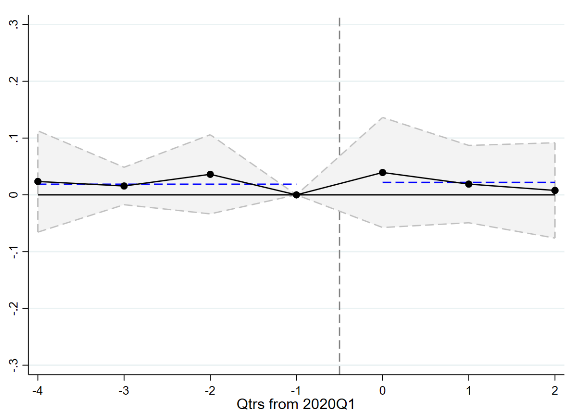

NOTE.——This figure plots coefficient estimates from Equation 5. The characteristic is either (a) the natural log of the income, (b) a borrower's income is verified, or (c) a borrower's employment is verified. The x-axis corresponds to the number of quarters from the treatment date. The quarter $\tau = -1$ is the reference quarter. The circles correspond to the coefficient estimates, and the vertical bars correspond to 95 percent confidence intervals. Standard errors are clustered at the lender-credit score bin level.

Figure A.1: Discontinuities in probability of automated underwriting

(a) Nissan

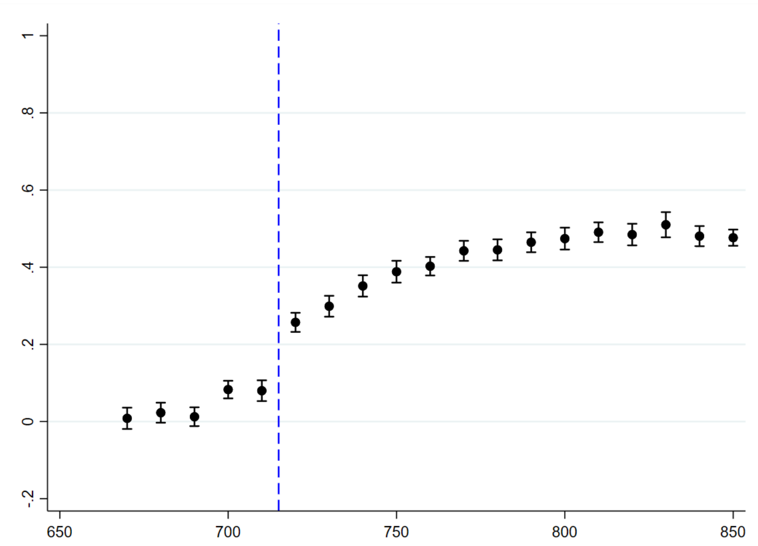

(b) Volkswagen

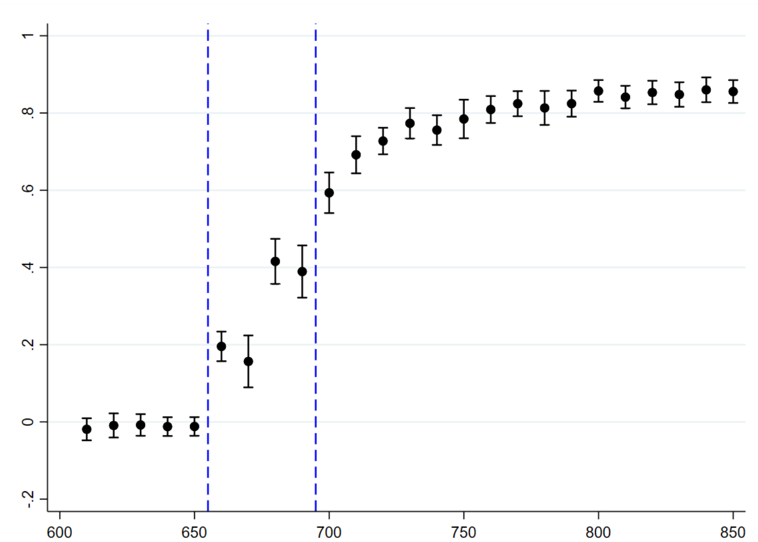

NOTE.——This figure plots the conditional average of automated underwriting across credit score bins for two lenders: (a) Nissan, and (b) Volkswagen. The circles correspond to the coefficient estimates, and the vertical bars correspond to 95 percent confidence intervals. Standard errors are clustered at the lender level.