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

Sarah Khoei†
University of Kentucky

September 10, 2025

(Job Market Paper)

Abstract

While prior studies find that automated underwriting outperforms manual underwriting, I show that there is significant heterogeneity in the adoption of automated underwriting both within and across lenders. 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 combination of difference-in-differences and regression discontinuity designs, 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 the pandemic period. The effect is particularly pronounced among higher-risk segments of borrowers, whose income and employment were more likely to be disrupted by the pandemic. Together, these results highlight 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 the auto lending industry.

Keywords: Automated underwriting; Auto loans; COVID-19 pandemic; FinTech.

JEL classification: G21; G28; G51; D81; O33

*I am deeply grateful to my committee members, David Sovich, Russell Jame, and Taylor Begley, for their invaluable insights and guidance throughout my Ph.D. program. I also extend my gratitude to the seminar participants at the University of Kentucky for their constructive feedback.

†Contact: sarah.khoei@uky.edu.

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 reflects the limitations of automated systems during periods of economic uncertainty, when their relative performance deteriorates. While a growing body of literature has examined the role of technology in household consumer markets, the focus has been on stable periods, leaving the comparative performance of technology during unexpected shocks largely unexplored. In this paper, I examine the performance of automated versus human underwriting in the face of such 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. This adaptability becomes particularly valuable when economic fundamentals shift rapidly, as soft information and real-time judgment can capture risks that historical models miss.

To test this hypothesis, I use the U.S. auto loan market as an ideal empirical setting and a difference-in-differences design to evaluate the performance of automated and manual underwriting before and after the COVID-19 pandemic, a significant unexpected shock. The U.S. auto loan market, the second-largest segment of consumer credit with $1.65 trillion in outstanding debt as of Q2 2025 (Federal Reserve Bank of New York 2025), 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, covering loans originated after 2017, aligns with this timeline and supports 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 cleanly capture shifts in their relative performance under economic stress. Absent such random assignment, 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 embed this fuzzy RDD into a difference-in-differences framework to compare performance across stable and unstable economic conditions. Fundamentally, this design compares loans just above the FICO thresholds, where automation is more prevalent, with those just below, where human underwriting is more common, across the pre- and post-pandemic periods. The results provide compelling evidence of the brittleness of automated systems in the face of economic shocks. During stable economic conditions, I find no significant performance differences between the two methods. This result is consistent with Jansen, Nguyen, and Shams 2024, who show that the overall advantage of automation is concentrated among borrowers with lower credit scores, whereas at higher credit scores, where the discontinuities in this setting occur, performance is essentially the same. After the onset of the pandemic, however, this pattern changes sharply. Default rates on automatically underwritten loans increase relative to those underwritten by humans, and manual underwriting outperforms automated systems. This shift 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, were more adaptable to the changing conditions.

Why does the performance of automated underwriting deteriorate after the onset of the COVID-19 pandemic? To investigate this question, I first examine whether changes in borrower composition could account for the observed differences in loan outcomes between automated and manual underwriting. The analysis shows no meaningful shifts in borrower characteristics such as income, employment and income verification, or co-obligor status across the two groups after the pandemic. I then turn to loan contract terms and find that maturities, loan-to-value ratios, and loan amounts increased for automated loans relative to those underwritten manually. Since automated models cannot adapt to new conditions, which is consistent with the findings of Ben-David, J. Johnson, and Stulz 2025 who show that FinTech lenders continued to issue loans on essentially unchanged terms during pandemic despite dramatic changes in the economic environment, the observed differences in contract terms reflect adjustments by human underwriters. In response to heightened uncertainty, they tightened standards by reducing loan-to-value ratios, shortening maturities, and approving smaller loans in order to limit risk exposure. Human underwriters demonstrated superior ability to incorporate forward-looking information, for example by discounting inflated credit scores that overstated borrower strength during the forbearance period and anticipating layoffs in industries directly affected by the crisis. In addition, the analysis shows that interest rates on automated loans increased relative to those on manually underwritten loans. This pattern suggests that lenders recognized the heightened risk associated with automated models and sought to compensate by charging higher rates on loans originated through automation. Such pricing adjustments reflect an institutional effort to offset risk in the absence of timely model recalibration, which in turn can partly explain the higher default rates observed for automated loans.

While the within-lender discontinuities provide credible causal evidence, they are inherently local, capturing comparisons only among borrowers near specific FICO thresholds. Since these thresholds are concentrated at relatively high credit scores, the design also provides limited scope for cross-sectional analysis, particularly for riskier borrower segments. To broaden the analysis, I complement this approach with an across-lender design that compares the portfolio performance of fully automated lenders with that of lenders combining manual and automated underwriting. Consistent with the within-lender results, my findings indicate that fully automated lenders experienced a deterioration in loan outcomes relative to mixed lenders during the post-pandemic period. This portfolio-level analysis confirms that the local effects identified in the RDD framework extend to broader patterns across the credit distribution.

Next, I examine how the effects vary across different types of borrowers. Cross-sectional analysis shows that the rise in defaults at fully automated lenders is concentrated among borrowers with below-median income, below-median credit scores, and above-median loan-to-value ratios. This finding is consistent with prior research documenting that these borrowers are disproportionately vulnerable to unexpected shocks, making models trained on pre-pandemic data especially prone to misprediction in these groups. Taken together, the evidence suggests that human underwriters add the greatest value when screening higher-risk borrowers, precisely where traditional hard information became less predictive during the pandemic. In this environment, the capacity of human underwriters to exercise judgment and incorporate soft information, such as anticipating layoffs in industries heavily affected by the crisis, proved especially important. Consistent with Iyer et al. 2016, who show that soft information plays a greater role in the evaluation of low-credit-score borrowers, my results underscore that discretion became particularly valuable for these segments during the pandemic, enabling mixed lenders to manage risk more effectively than their fully automated counterparts. Moreover, my results reveal that the increase in interest rates on automated loans was likewise concentrated among these same low-income and low-credit-score borrowers, indicating that lenders raised prices precisely in the segments where their models were most vulnerable to breakdown.

An important question is whether lenders that relied on a combination of human and automated underwriting responded to heightened uncertainty by shifting a greater share of loans to manual review in order to offset the weaknesses of automated systems. My results indicate that such lenders did, in fact, increase the proportion of loans originated through human underwriters. Yet the scale of this adjustment remained modest. Two constraints help explain this limited shift. First, existing underwriters faced capacity limitations and could not easily expand their review workload, thereby restricting the extent of reallocation. Second, expanding capacity by hiring additional staff was both time-consuming and resource-intensive, and during the pandemic these frictions were especially severe. As a result, lenders' ability to substitute toward discretionary review was sharply constrained. Similar bottlenecks are documented by Fuster et al. 2021, who show that traditional mortgage lenders struggled to expand loan officer capacity during periods of heightened demand in the COVID-19 pandemic.

The validity of the difference-in-differences analysis rests on the assumption that, absent the pandemic, the performance of loans originated by human underwriters and those originated through automated systems would have evolved along comparable trajectories. To evaluate this assumption, I examine pre-pandemic default rates across the two underwriting methods. The results show no significant divergence in trends between the groups, indicating that automated and human underwriting followed parallel paths prior to the COVID-19 shock. This evidence 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 performance, 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 underwriting 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 examines cross-sectional heterogeneity, and 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

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

proposed loan terms.

Upon receiving a credit application, the lender generally obtains a credit report on the 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, they independently evaluates the application based on the lender's established underwriting guidelines and professional judgment. 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. Unlike automated systems, human underwriters are able to incorporate qualitative information, exercise judgment, and adapt to circumstances not captured by quantitative models. This discretion allows them to interpret the broader context of an applicant's financial situation, assess the plausibility of reported information, and adjust credit standards in response to changing conditions. For example, they may recognize that a borrower's credit score remains depressed due to a past bankruptcy, which typically lingers on credit reports for seven to ten years, even when the borrower's current financial standing is strong. They can also identify potential inconsistencies, such as implausibly high stated income relative to occupation, and request further documentation. In addition, human underwriters can respond to broader economic conditions, tightening standards when particular industries experience widespread layoffs or loosening them when risks abate.

After evaluating the strengths and weaknesses of an application, the credit analyst decides whether to approve or reject it. 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]

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

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 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 document 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 displays striking variation in the adoption of underwriting methods, with human judgment continuing to play an important role despite the efficiency gains associated with automation. Figure 2 highlights these differences across lenders. A group of institutions relies almost entirely on automated systems, including Capital One, CarMax, Ford Credit, Exeter Finance, Hyundai Capital, Mercedes-Benz Financial, and World Omni. In contrast, several lenders follow a hybrid approach that integrates both automated and manual underwriting, such as Ally, American Honda Finance, BMW Financial Services,

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).

Nissan Finance, GM Financial, Santander, Toyota Financial Services, and Volkswagen Financial Services. Fifth Third Bank stands out as the only lender in the sample to depend exclusively on manual underwriting. An especially noteworthy feature of this variation is that even captive finance companies serving similar borrower bases and specializing in new car sales differ in their reliance on human discretion. For example, Nissan Finance continues to underwrite more than half of its loans manually, whereas GM Financial relies far more heavily on automation. This heterogeneity illustrates that automation has not displaced human review uniformly across the industry and underscores the enduring value of discretion even in markets where technological tools are widely available.

The heterogeneity in underwriting practices across lenders is not driven by changes over time. Figure 3 shows that the share of loans processed through automated systems remained remarkably stable between 2018 and 2022. Even after the onset of the COVID-19 pandemic, there is little evidence of systematic shifts toward or away from automation. This stability is particularly striking given both the rapid improvements in technology and the severe disruptions created by the pandemic. Taken together, the evidence suggests that the cross-sectional variation in automation reflects persistent institutional choices, highlighting the importance of understanding why some lenders continue to preserve a role for human discretion while others rely almost exclusively on automated systems.

Several institutional and market-level factors help explain the cross-lender heterogeneity in automation. At its core, underwriting strategy reflects a cost–benefit tradeoff. Automation offers clear advantages in speed, scalability, and operational efficiency, enabling lenders to process large volumes of applications quickly and at lower cost. FinTechs and platform-based lenders, in particular, emphasize these benefits and are therefore more inclined to adopt fully automated systems. In contrast, traditional banks and manufacturer-affiliated captives often place greater weight on risk management and preserve manual review, accepting higher costs and slower processing in exchange for greater flexibility and resilience. For some lenders, especially those closely tied to dealerships, retaining human underwriters also facilitates the incorporation of dealer input and the handling of complex, bundled products such as financing, insurance, and service add-ons. These arrangements frequently involve borrower- or deal-specific soft information that is difficult to codify in automated models, further reinforcing the value of human discretion.

In addition to these explanations, an important reason some lenders preserve manual underwriting capacity is the belief that human underwriters are better equipped to navigate periods of macroeconomic uncertainty. When the historical patterns embedded in automated models no longer hold, human underwriters can exercise judgment, adapt criteria, and incorporate evolving information in ways that rigid models cannot. This paper investigates whether such advantages translate into observable differences in loan performance across underwriting methods during the COVID-19 pandemic, thereby providing new evidence on the conditions under which human judgment adds value.

To test the hypothesis that lenders continue to rely on human underwriters because automated systems perform worse during periods of economic uncertainty, I restrict the estimation sample to auto loans originated in the two years before and after the onset of the COVID-19 pandemic (January 2018 through December 2022). This window enables a direct comparison of the performance of automated and manual underwriting systems under both stable and uncertain economic conditions. I further limit the sample to loans with complete information on key contract and borrower characteristics, including interest rate, loan amount, maturity, scheduled monthly payment, vehicle condition (new or used), make, model, model year, vehicle value, borrower credit score, and income. To ensure data reliability, I exclude loans with interest rates above 30 percent and credit scores below 500. I also restrict the sample to new vehicles, since the dataset does not contain mileage information that would allow for controlling the quality of used cars. In addition, I keep only lenders that originate loans in both the pre- and post-COVID periods, and I exclude Harley-Davidson, which primarily finances motorcycles rather than automobiles. Finally, I winsorize vehicle value, borrower income, loan-to-value ratio, and payment-to-income ratio at the 1 percent level to reduce the influence of outliers.

Table 1 presents descriptive statistics for loans at the time of origination. The average loan in the sample has an interest rate of 4.5 percent, a scheduled monthly payment of $535, a vehicle price of $33,940, a maturity of 68 months, and a loan amount of $31,281. The average loan to value ratio is 94 percent. The borrowers in the sample have an average credit score of 747 and an annual household income of $99,808. The unconditional default rates are 1.2 percent within 12 months, 2.9 percent within 24 months, and 4.7 percent within 36 months of origination.

The two right-most columns of Table 1 compare loans originated by fully automated lenders with those issued by partially automated lenders, restricting the sample to loans originated prior to the pandemic. Several differences emerge between the two groups. Loans from fully automated lenders feature higher loan amountss ($30,837 versus $29,638), longer maturities (69 months versus 67 months), lower interest rates (3.6 percent versus 5.8 percent), and higher loan-to-value ratios (0.95 versus 0.93). Borrowers served by fully automated lenders also have higher average credit scores (757 versus 735) but somewhat lower household incomes ($94,596 versus $98,535) compared to those borrowing from partially automated lenders.

To mitigate concerns about endogeneity in lenders' assignment of automated versus manual underwriting, I focus on lenders with identifiable discontinuities in the probability of automation. These discontinuities generate quasi-random variation in the likelihood of automated underwriting, which enables an evaluation of how the relative performance of automated and manual systems evolved before and after the onset of the COVID-19 pandemic. Using the data-driven procedures described in Section 4.1.1, I detect such discontinuities across credit score thresholds for two captive lenders, Nissan Finance and Volkswagen Financial Services. The final estimation sample consists of 179,968 auto loans originated by these lenders in the neighborhoods of the identified cutoffs. Table 2 reports descriptive statistics for these loans. On average, loans in this restricted sample carry an interest rate of 3.3 percent, a scheduled monthly payment of $495, a vehicle price of $29,818, a maturity of 69 months, and a loan amount of $30,000. The mean loan-to-value ratio is 101.8 percent. Borrowers have an average credit score of 722 and an annual household income of $84,453. The unconditional default rates are 0.3 percent within 12 months, 1.4 percent within 24 months, and 2.7 percent within 36 months of origination.

The two right-most columns of Table 2 compare loans originated just above the FICO score cutoffs with those originated just below them within the restricted sample. The comparison is limited to loans issued prior to the pandemic. Several differences emerge between the two groups. Loans originated above the cutoffs carry lower interest rates (4.0 percent versus 4.5 percent) and are associated with higher vehicle values ($29,332 versus $28,811). In addition, borrowers with loans above the cutoffs have higher average credit scores (728 versus 711) and higher incomes ($84,430 versus $82,418) compared to those with loans below the cutoffs.

Although these loans show observable time-invariant differences, the difference-in-differences model in Section 4.1.3 addresses these differences by including fixed effects for lender, vehicle, and borrower characteristics. As shown in Figures 7 and A.3, 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

This section outlines the empirical methodology. I begin with a loan-level analysis that compares the performance of automated and manually underwritten loans, exploiting quasi-random within-lender variation in automation probability around lender-specific FICO score thresholds. This design is supported by tests of standard identification assumptions. In a complementary analysis, I extend the focus to the lender level, comparing the portfolio performance of fully automated lenders with that of mixed lenders. Taken together, these approaches provide a comprehensive evaluation of the relative performance of automated and manual underwriting in the period surrounding the COVID-19 pandemic.

4.1 Quasi-Random Loan-Level Analysis

To address the endogeneity in the assignment of automated versus manual underwriting, I exploit quasi-random variation in the probability of automation generated by lender-specific discontinuities at various FICO score thresholds. These discontinuities create sharp changes in the likelihood of automated underwriting for otherwise comparable borrowers. By focusing on loans originated just above and below the identified thresholds, I compare loans with similar observable characteristics but different probabilities of automation, thereby mitigating concerns about selection bias.

4.1.1 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 500, with the first bin covering scores from 500 to 509, the second from 510 to 519, and so forth, up to scores of 900. 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 500 to 509).

This procedure identifies discontinuities in the likelihood of automation for two lenders: Nissan Finance 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. In Figure A.1, Panel (a) highlights a discrete jump of 16 percentage points in the likelihood of automation at a FICO score of 720 for Nissan. Panel (b) illustrates similar jumps for Volkswagen at thresholds of 660 and 700, with magnitudes of 20 and

24 percentage points, respectively. Because these estimates are conditioned on observable borrower and loan characteristics, they suggest that the discontinuities are not driven by shifts in composition and can be interpreted as quasi-exogenous variation in the probability of automation. In addition, Figure 4 plots the unconditional average automation rate in 2-point FICO bins. The patterns closely mirror the regression results, reinforcing the robustness of the discontinuities.

4.1.2 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 1 and 2:

$$\text{Automated}_{i,t} = \omega_1 + \Gamma_1 \cdot T_i + f(x_i) + T_i \cdot g(x_i) + \delta_{s,t} + \delta_{l,t} + \delta_{I,t} + \delta_{v,t} + \nu_{i,t}, \tag{1}$$

$$\text{Automated}\times\text{After}_{i,t} = \omega_2 + \Gamma_2 \cdot T_i \cdot \text{After}_t + f(x_i) + T_i \cdot g(x_i) + \delta_{s,t} + \delta_{l,t} + \delta_{I,t} + \delta_{v,t} + \zeta_{i,t}, \tag{2}$$

where $\text{Automated}_{i,t}$ is a binary indicator equal to 1 if loan i originated in quarter t is underwritten through 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 After_t is a dummy variable equal to 1 for loans originated after the onset of the COVID-19 pandemic and 0 otherwise. The variable $\text{Automated}\times\text{After}_{i,t}$ captures the combined effect of automation and the post-pandemic period. 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. This specification includes fixed effects to control for lender ($\delta_{l,t}$), state ($\delta_{s,t}$), vehicle ($\delta_{v,t}$), and borrowers' income ($\delta_{I,t}$). The estimation sample is restricted to loans with credit scores within 20 points of an automated underwriting discontinuity.

Table 3 reports the first-stage estimates from Equations 1 and 2. Panel A shows results where the dependent variable is an indicator equal to one if a loan was originated through automated underwriting. Across all specifications, the coefficient on the treatment indicator, T_i, is positive and highly statistically significant, ranging between 0.084 and 0.085. This implies that crossing the FICO score threshold increases the probability of a loan being originated through automated underwriting by roughly 8.4 to 8.5 percentage points. Panel B shifts the focus to the interaction term Automated Underwriting $\times$ After as the dependent variable, thereby capturing variation in automation specifically in the post-treatment period. Here, the coefficient on the treatment-after interaction is consistently large, positive, and highly statistically significant, with magnitudes around 0.055 to 0.056. This indicates that the probability of automated underwriting increases by 5.5 to 5.6 percentage points after the onset of the pandemic at the FICO score thresholds. Again, the results are robust across specifications with different sets of fixed effects. Taken together, Panels A and B demonstrate a strong and robust first stage: the identified thresholds create substantial quasi-random variation in automation both overall and conditional on the post-pandemic period. These results provide confidence in the strength of the instruments, with F-statistics far exceeding conventional thresholds, thereby supporting the validity of the research design.

4.1.3 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:

$$y_{i,t} = \alpha + \gamma \cdot \widehat{\text{Automated}}_{i,t} + \beta \cdot \widehat{\text{Automated}\times}\text{After}_{i,t} + f(x_i) + T_i \cdot g(x_i) + \delta_{s,t} + \delta_{l,t} + \delta_{I,t} + \delta_{v,t} + \eta_{i,t}, \tag{3}$$

where the outcome variable is an indicator equal to one if loan i originated in quarter t defaults within 12, 24, or 36 months of origination. $\widehat{\text{Automated}}_{i,t}$ and $\widehat{\text{Automated}\times}\text{After}_{i,t}$ are the predicted values from the first-stage regressions (Equations 1 and 2). The instrument T_i is a dummy variable equal to 1 if the credit score is above the threshold. 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 FICO score (running variable) 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 based on auto loans originated by Nissan and Volkswagen between 2018 and 2022, restricting the sample to loans issued just above and below the identified FICO score thresholds.

Table 4 presents the coefficient estimates from Equation 3. Across all specifications, the coefficient of interest, β, is positive and statistically significant, indicating a substantial increase in default rates for loans underwritten automatically compared to those underwritten manually in the post-COVID period. Specifically, automated loans originated after the onset of the pandemic are 3.6 percentage points more likely to default within 12 months, 9.1 percentage points more likely to default within 24 months, and 14.3 percentage points more likely to default within 36 months relative to manually underwritten loans around the same cutoffs. Furthermore, by contrast, the coefficient on $\widehat{\text{Automated}}_{i,t}$,$\gamma$, captures the pre-pandemic difference in default rates between automated and manual underwriting. This coefficient is small and statistically insignificant, indicating that there were no systematic differences in default rates between automated and manual underwriting in the pre-pandemic period. This finding aligns with Jansen, Nguyen, and Shams 2024, who show that the performance advantage of automation is concentrated among borrowers with lower credit scores. At higher credit scores, where the discontinuities in this setting occur, automated and manual underwriting perform similarly. The running variable and its interactions are generally insignificant, suggesting no evidence of sorting or functional form sensitivity around the thresholds. Taken together, these results demonstrate that the relative performance of automated underwriting deteriorated significantly after the onset of the COVID-19 pandemic. While automated and manual underwriting performed similarly in stable conditions, automated systems experienced substantially higher default rates when confronted with heightened economic uncertainty.

To analyze how the impact of COVID-19 evolved over the sample period, I estimate Equation 4:

$$y_{i,t} = \alpha + \sum_{\tau=-4}^{3} \Gamma_\tau \cdot \widehat{\text{Automated}}_i \cdot \text{D}_\tau + \delta_a + \delta_{s,t} + \delta_{l,t} + \delta_{I,t} + \delta_{v,t} + \varepsilon_{i,t}, \tag{4}$$

where D_τ is equal to one if period t is τ half-year intervals from the treatment date. I exclude the period prior to the treatment date, defined as the onset of COVID-19, ($\tau = -1$) as the reference category. Therefore, the Γ_τ coefficient captures the average difference in default rate between loans originated by automated systems and those underwritten manually in half-year τ relative to the average difference observed in default rate six months prior to the treatment date. I also control for the time-invariant variation in the likelihood of automation by adding δ_a. The results, presented in Figure 7, panels (a)–(c), show that default rates for automatically underwritten loans increased substantially in the post-COVID period. 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. Figure A.2 panels (a)–(c) in the Appendix extends the analysis to three years beyond the onset of the pandemic. The results reveal a reversal in the default gap: the difference in performance between automated and manual underwriting narrows, and default rates on automated loans decline relative to the immediate post-COVID period. This pattern suggests that automated systems were gradually updated and recalibrated as new data reflecting post-pandemic conditions was incorporated, leading to a recovery in model performance. In addition to the semiannual specification, Figures A.3 and A.4 present the dynamic effects estimated at quarterly intervals. The results closely mirror the semiannual analysis, showing a sharp deterioration in the performance of automated underwriting in the immediate post-COVID period, followed by a gradual reversal as models were updated with post-pandemic data. The consistency across both temporal aggregations reinforces the robustness of the findings.

4.1.4 Quasi-Random Variation: Identification Assumptions

In addition to the relevance condition (shown in Table 3 as well as Figures 4 and A.1), the internal validity of the fuzzy regression discontinuity design (RDD) 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 implies that no other loan contract terms exhibit discontinuous changes around the credit score thresholds. To test this assumption, I estimate Equation 5:

$$y_{i,t} = \alpha + \gamma \cdot \widehat{\text{Automated}}_{i,t} + \beta \cdot \widehat{\text{Automated} \times \text{After}}_{i,t} + f(x_i) + T_i \cdot g(x_i) + \delta_{s,t} + \delta_{l,t} + \delta_{I,t} + \delta_{v,t} + \eta_{i,t}, \quad (5)$$

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 instrument T_i is a dummy variable equal to 1 if the credit score is above the threshold. 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. This specification includes fixed effects to control for lender ($\delta_{l,t}$), state ($\delta_{s,t}$), vehicle ($\delta_{v,t}$), and borrowers' income ($\delta_{I,t}$).

Table 5 reports formal regression tests of whether loan terms change around the credit score thresholds. The estimated coefficients, γ, for interest rate, loan maturity, loan-to-value ratio, and loan amount are uniformly small and statistically insignificant. This indicates that, in the pre-treatment period, assignment to automated versus manual underwriting was not associated with mechanically different contract terms. These findings support the validity of the exclusion restriction, showing that the discontinuities do not mechanically alter loan terms but instead operate only through their effect on the likelihood of automation.

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 5, using borrower income, the presence of a co-obligor, and indicators for whether income or employment is verified as outcome variables. The results, reported in Table 6, show no evidence of systematic changes in these characteristics around the thresholds. The estimated coefficients, γ, are neither economically nor statistically distinguishable from zero. Figures 5 and 6 further illustrate the absence of discontinuities in these variables for Nissan and Volkswagen, providing robust support for the validity of the local continuity assumption.

4.1.5 Why did automated underwriting struggle during the pandemic shock?

To understand why the performance of automated underwriting deteriorated after the pandemic, I begin by examining whether borrower composition changed during this period. A shift in the pool of borrowers could explain the performance gap if automated and manual underwriting were applied to systematically different groups after COVID-19. Table 6 presents the results. The coefficients on Automated $\times$ After are small and statistically insignificant, indicating no meaningful change in borrower characteristics across underwriting methods in the post-pandemic period.

Next, I turn to loan contract terms. Table 5 shows that loans originated through automated systems were issued with longer maturities, higher loan-to-value ratios, and larger principal balances compared to those underwritten manually. These patterns suggest that automated underwriting continued to extend riskier credit after the shock. Human underwriters, by contrast, incorporated real-time information and tightened standards accordingly. They recognized that credit scores were temporarily inflated by government forbearance programs and stimulus transfers[3], and they anticipated elevated layoff risks among borrowers in industries disproportionately affected by COVID-19. By discounting these distortions, human underwriters originated loans on more conservative terms. Automated systems, however, could not adjust in the short run, as retraining requires updated data that become available with a lag. Consequently, automated underwriting continued to operate as if conditions had not changed, systematically approving riskier contracts in the post-pandemic period. In addition, the coefficient on Automated $\times$ After for interest rates indicates that, following the pandemic, automated loans carried higher rates relative to manually under-

3. During the COVID-19 pandemic, credit scores were temporarily and artificially elevated due to government forbearance programs and stimulus transfers. For more discussion, see MarketWatch [link1] and CFPB [link2].

written loans. This pattern suggests that lenders recognized the elevated risks associated with automated models and sought to compensate by increasing the cost of credit on loans originated through automation. Such pricing adjustments reflect an institutional attempt to offset risk in the absence of timely model recalibration and may help explain, at least in part, the higher default rates observed for automated loans.

A natural question is why lenders were unable to promptly update their lending standards after the onset of COVID-19. Several explanations are plausible. First, recalibrating automated underwriting systems requires sufficient post-pandemic repayment data, which can only be observed with a lag. Without reliable performance information, lenders could not immediately adjust model parameters or retrain their algorithms. Second, even when lenders sought to modify their models, any major change to underwriting rules typically required regulatory approval from the Consumer Financial Protection Bureau (CFPB). This oversight ensures that adjustments do not inadvertently introduce or reinforce discriminatory lending practices. Because modifications to underwriting criteria may correlate with borrower characteristics protected under fair lending laws, such as race or gender, the approval process can be lengthy. These delays were likely compounded by the operational disruptions of the pandemic, including the widespread shift to remote regulatory procedures.

4.2 Lender-Level Analysis

While the loan-level design exploits within-lender discontinuities to identify quasi-random variation in automation, it remains local in nature and restricted to borrowers near specific FICO score thresholds. Because these thresholds are concentrated at relatively high credit scores, the design provides limited leverage for cross-sectional analysis, particularly among riskier segments of the market. To complement this local evidence, I therefore turn to an across-lender design that compares portfolio outcomes for fully automated lenders with those for lenders that combine manual and automated underwriting. This approach expands the scope of inference and allows cross-sectional analyses that are not feasible within the near-threshold framework. Formally, I estimate the following regression:

$$y_{i,t} = \alpha + \gamma \cdot \text{Automated lender}_i + \beta \cdot \text{Automated lender} \times \text{After}_{i,t} + \delta_{s,t} + \delta_{l,t} + \delta_{c,t} + \delta_{I,t} + \delta_{v,t} + \eta_{i,t}, \tag{6}$$

where the outcome variable is an indicator for whether loan i originated in quarter t defaults within 12, 24, or 36 months. Automated lender$_i$ is an indicator equal to one if the originating lender is fully automated and zero if the lender employs a mixed approach combining manual and automated underwriting. This specification includes a rich set of 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, measures the differential change in default rates after the onset of COVID-19 for loans originated by fully automated lenders relative to those originated by partially automated lenders.

Table 7, Panel A, reports the lender-level difference-in-differences estimates of loan performances. The coefficients on Automated Lender $\times$ After are positive and statistically significant across all horizons, indicating that defaults increased disproportionately at lenders that rely exclusively on automated underwriting compared to those that combine manual and automated methods. Specifically, default rates at fully automated lenders rose by 1.2 percentage points at 12 months, 1.9 percentage points at 24 months, and 2.6 percentage points at 36 months relative to mixed lenders. These findings reinforce the loan-level results by showing that the deterioration of automated underwriting performance during the pandemic was not confined to local discontinuities around credit score thresholds but was also evident in the aggregate portfolios of fully automated lenders. The results highlight that lenders relying exclusively on automation faced systematically higher default risk when economic conditions shifted abruptly, whereas mixed lenders were better able to adapt through discretionary review.

Consistent with the loan-level analysis, Table 7, Panel A, also shows that fully automated lenders passed on part of the elevated model risk to consumers by charging higher interest rates after the onset of the pandemic. On average, the portfolio-level interest rate at these lenders increased by 0.004, or roughly 40 basis points, relative to lenders that relied on a combination of automated and manual underwriting. This finding suggests that, in the absence of timely model recalibration, lenders resorted to pricing adjustments to partially offset the heightened default risk associated with automated systems. By contrast, Panel B shows no evidence of changes in borrower composition following the pandemic. The estimated coefficients on the presence of a co-obligor, income verification, employment verification, and natural log of income are all small and statistically insignificant, reinforcing the conclusion that the deterioration in automated underwriting performance was not driven by shifts in the underlying borrower pool.

4.2.1 Did lenders expand manual underwriting after the pandemic?

An important question is whether lenders increased their reliance on manual underwriting after the onset of the pandemic, recognizing that automated models were more prone to error under rapidly changing conditions. To address this, I estimate Equation 6 using the share of loans originated through automation as the dependent variable. As shown in Table 8, the coefficient estimates are negative and statistically significant across all specifications. On average, the share of automated underwriting declined by 2.6 percentage points, indicating that lenders shifted a portion of their originations toward manual review. The magnitude of this adjustment, however, was limited. Two constraints help explain the restricted shift toward manual underwriting. First, existing underwriters faced strict capacity limits and could not easily expand their review workload, which sharply constrained reallocation. Second, expanding capacity by hiring new staff was both time-consuming and resource-intensive. The auto loan market has chronic staffing shortages, and training new underwriters requires substantial time to develop judgment skills, while even experienced staff need time to adapt to a lender's policies, dealer relationships, and risk models.[4] These frictions were particularly severe during the pandemic, when disruptions to labor markets and organizational operations further delayed the ability of lenders to expand discretionary review. Consistent evidence is provided by Fuster et al. (2021), who show that traditional mortgage lenders struggled to expand loan officer capacity during periods of heightened demand in the COVID-19 pandemic. It is worth noting that fully automated lenders could not reallocate loans to manual underwriters, as they had invested entirely in automation and lacked the infrastructure to support discretionary review.

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 along the median of borrowers' income, credit score, and loan-to-value ratio. To capture the differential impacts of the pandemic across these income groups, I estimate the following triple-differences (DDD) model using Equation 7:

$$\begin{aligned} y_{i,t} = \alpha + \beta \cdot \mathrm{Z}_{ik} \cdot \text{Automated}\widehat{\text{ lender}} \text{ X After}_{i,t} + \Gamma \cdot \text{Automated}\widehat{\text{ lender}}\text{XAfter}_{i,t} + \\ \theta \cdot \mathrm{Z}_{ik} \cdot \widehat{\text{Automated lender}}_{i,t} + \delta_{s,t} + \delta_{c,t} + \delta_{I,t} + \delta_{v,t} + \varepsilon_{i,t} \end{aligned} \tag{7}$$

where the outcome variable is an indicator for whether loan i originated in quarter t defaults within 12, 24, or 36 months. The variable K refers to the borrower characteristic of interest, which can be Low Income, Low Credit Score, or High LTV. Z_{ik} is an indicator for whether borrower i belongs to subgroup k. The coefficient of interest, β, measures how the effect of the COVID-19 pandemic on loan performance differs between automated and

4. These staffing challenges are increasingly evident in auto finance industry—see Digital Dealer[link3].

manual underwriting specifically for borrowers in the lowest income group, the lowest credit score group, or the highest loan-to-value group compared with their respective benchmark groups.

Table 9, Panel A, presents evidence on cross-sectional heterogeneity in the lender-level effects of automation after the pandemic shock. The results show that the increase in default rates at fully automated lenders is concentrated among the riskiest borrower segments—those with the lowest income, the lowest credit scores, and the highest loan-to-value ratios—who are most vulnerable to economic shocks. Specifically, 12-month default rates are 80 basis points higher for low-income borrowers, 2.6 percentage points higher for low-credit-score borrowers, and 50 basis points higher for high-LTV borrowers. All of these estimates are statistically significant. These findings suggest that automated models performed particularly poorly in predicting repayment behavior for borrowers with weaker financial profiles, underscoring their limitations when faced with heightened uncertainty.

Next, I estimate Equation 7 using interest rate as the dependent variable. Consistent with the heterogeneous effects observed for defaults, Table 9, Panel B, shows that lenders increased rates more sharply for low-income, low-credit-score, and high-LTV borrowers. This pattern indicates that lenders raised the cost of credit precisely in those segments where automated models exhibited the greatest misprediction, reflecting an attempt to price compensate for elevated risk.

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 puzzle 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 and lower-credit-score 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.

Federal Reserve Bank of New York. 2025. *New York Fed to Release Q2 2025 Household Debt and Credit Report on August 5, 2025.* Media advisory. Accessed: 2025-09-02. https://www.newyorkfed.org/newsevents/mediaadvisory/2025/0731-2025.

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.

Iyer, Rajkamal, Asim Ijaz Khwaja, Erzo FP Luttmer, and Kelly Shue. 2016. "Screening peers softly: Inferring the quality of small borrowers." *Management Science* 62 (6): 1554–1577.

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 (Full sample)

	(1) Mean	(2) SD	(3) P10	(4) P25	(5) P50	(6) P75	(7) P90	(8) Fully automated lender	(9) Partially automated lender
Automated underwriting	0.837	0.369	0	1	1	1	1	1	0.756
Loan amount	31,281	11,987	18,334	23,013	29,289	37,336	46,561	30,837	29,638
Interest rate	0.045	0.047	0.000	0.019	0.030	0.056	0.108	0.036	0.058
Monthly payment	535	201	320	398	502	634	785	508	533
Maturity (months)	68	9	61	61	73	74	76	69	67
Loan-to-Value	0.941	0.235	0.619	0.792	0.957	1.101	1.233	0.948	0.929
Car value	33,940	11,745	21,721	25,435	30,822	40,433	49,246	33,217	32,734
Credit score	747	88	615	694	759	818	851	757	735
Income	99,808	67,030	42,005	55,431	81,042	120,000	180,000	94,596	98,535
Co-Obligor	0.330	0.470	0	0	0	1	1	0.330	0.338
12-month default	0.012	0.110	0	0	0	0	0	0.004	0.025
24-month default	0.029	0.169	0	0	0	0	0	0.013	0.052
36-month default	0.047	0.212	0	0	0	0	0	0.022	0.083
Observations	4,452,286							705,496	1,616,283

NOTE.—This table summarizes the full sample of 4,452,286 auto loans originated in the two years before and after the onset of COVID-19. Descriptive statistics are as of the loan origination date. In Columns 8 and 9, I compare auto loans originated by fully automated lenders to loans originated by partially automated lenders during the pre-treatment period.

Table 2: Descriptive statistics (around identified cutoffs)

	(1) Mean	(2) SD	(3) P10	(4) P25	(5) P50	(6) P75	(7) P90	(8) Above cutoffs	(9) Below cutoffs
Automated underwriting	0.337	0.473	0	0	0	1	1	0.328	0.212
Loan amount	29,999	10,587	18,861	23,057	28,254	34,890	42,850	29,407	29,464
Interest rate	0.033	0.025	0.000	0.009	0.029	0.051	0.067	0.040	0.045
Monthly payment	495	209	300	369	455	568	722	495	505
Maturity (months)	69	7	60	61	72	73	74	69	69
Loan-to-Value	1.018	0.212	0.737	0.889	1.031	1.169	1.276	1.018	1.037
Car value	29,818	9,657	20,360	23,711	27,583	33,322	41,986	29,332	28,811
Credit score	722	21	695	708	723	737	749	728	711
Income	84,453	54,692	36,037	48,082	70,045	102,000	150,000	84,430	82,418
Co-Obligor	0.262	0.440	0	0	0	1	1	0.281	0.274
12-month default	0.003	0.057	0	0	0	0	0	0.004	0.006
24-month default	0.014	0.116	0	0	0	0	0	0.018	0.025
36-month default	0.027	0.163	0	0	0	0	0	0.035	0.048
Observations	179,968							58,065	40,848

NOTE.— This table summarizes the sample of 179,968 auto loans originated around the identified cutoffs for Nissan and Volkswagen in the two years before and after the onset of COVID-19. Descriptive statistics are as of the loan origination date. In Columns 8 and 9, I compare auto loans above the identified cutoffs to loans below the identified cutoffs during the pre-treatment period.

Table 3: First-stage results

Panel A	(1) Automated	(2) Automated	(3) Automated	(4) Automated
Treat	0.085***	0.085***	0.084***	0.084***
	(8.69)	(8.38)	(8.59)	(8.78)
Running	0.000	0.000	0.000	0.000
	(0.38)	(0.23)	(0.25)	(0.30)
Running × Treat	0.003*	0.004**	0.004**	0.004**
	(1.91)	(2.01)	(2.18)	(2.12)
Running^2 × Treat	-0.000	-0.000	-0.000	-0.000
	(-0.58)	(-0.66)	(-0.80)	(-0.74)
Lender×Time FE	Yes	Yes	Yes	Yes
Vehicle×Time FE		Yes	Yes	Yes
State×Time FE			Yes	Yes
Income×Time FE				Yes
R^2	0.021	0.029	0.035	0.040
Obs	98,913	98,908	98,908	98,903

Panel B	(1) Automated X After	(2) Automated X After	(3) Automated X After	(4) Automated X After
Treat×After	0.056***	0.056***	0.055***	0.056***
	(9.95)	(10.31)	(10.40)	(10.25)
Running	0.000	0.000	0.000	0.000
	(1.27)	(1.21)	(1.19)	(1.23)
Running × Treat	-0.001	-0.001	-0.001	-0.001
	(-1.25)	(-1.14)	(-1.10)	(-1.17)
Running^2 × Treat	0.000**	0.000**	0.000**	0.000**
	(2.37)	(2.15)	(2.08)	(2.21)
Lender×Time FE	Yes	Yes	Yes	Yes
Vehicle×Time FE		Yes	Yes	Yes
State×Time FE			Yes	Yes
Income×Time FE				Yes
R^2	0.286	0.293	0.299	0.301
Obs	179,968	179,958	179,958	179,939

NOTE.— Panel A reports first-stage estimates from Equations 1. The dependent variable is an indicator equal to one if a loan was originated through automated underwriting. Panel B reports first-stage estimates from Equations 2. The dependent variable is the interaction term Automated× After. The sample is restricted to auto loans originated within two years before the treatment date. Standard errors are clustered at the running-variable level.
* Significant at the 10% level.
** Significant at the 5% level.
*** Significant at the 1% level.

Table 4: Second-stage results: loan performance around the cutoffs

	(1) 12-month default	(2) 24-month default	(3) 36-month default
Automated underwriting	0.013	0.005	0.016
	(0.53)	(0.12)	(0.24)
Automated × After	0.036*	0.091**	0.143**
	(1.66)	(2.32)	(2.37)
Running	-0.000	-0.000*	-0.001*
	(-1.29)	(-1.69)	(-1.80)
Running × Treat	-0.000	-0.000	-0.000
	(-0.30)	(-0.04)	(-0.13)
Running^2 × Treat	0.000	0.000	0.000
	(0.34)	(0.26)	(0.23)
Lender×Time FE	Yes	Yes	Yes
Vehicle×Time FE	Yes	Yes	Yes
State×Time FE	Yes	Yes	Yes
Income×Time FE	Yes	Yes	Yes
Obs	179,939	179,939	152,725

NOTE.—This table reports coefficient estimates from Equation 3. The dependent variable is an indicator equal to one if the borrower is in default 12, 24, or 36 months after loan origination. The sample is restricted to auto loans originated within two years before and two years after the treatment date. Standard errors are clustered at the running variable level.

* Significant at the 10% level.

** Significant at the 5% level.

*** Significant at the 1% level.

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

	(1) Interest rate	(2) log(maturity)	(3) Loan-to-Value	(4) log(loan amount)
Automated underwriting	-0.002	0.187	0.305	0.312
	(-0.11)	(1.22)	(0.99)	(1.05)
Automated × After	0.033*	0.169	0.566**	0.524*
	(1.72)	(1.20)	(2.01)	(1.91)
Running	-0.000*	-0.001	-0.002	-0.002
	(-1.85)	(-1.45)	(-1.47)	(-1.45)
Running × Treat	0.000	-0.000	-0.001	-0.001
	(0.16)	(-0.18)	(-0.28)	(-0.28)
Running^2 × Treat	-0.000	-0.000	0.000	0.000
	(-0.00)	(-0.00)	(0.13)	(0.13)
Lender×Time FE	Yes	Yes	Yes	Yes
Vehicle×Time FE	Yes	Yes	Yes	Yes
State×Time FE	Yes	Yes	Yes	Yes
Income×Time FE	Yes	Yes	Yes	Yes
Obs	179,939	179,939	179,939	179,939

NOTE.—This table reports coefficient estimates from Equations 3. The dependent variable is either the interest rate, the natural log of the maturity, loan to value ratio, or the natural log of the loan amount. Standard errors are clustered at the running variable 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) Co-obligor	(2) Income verification	(3) Employment verification	(4) log(income)
Automated underwriting	0.032	0.019	0.019	-0.195
	(0.41)	(0.97)	(0.97)	(-0.79)
Automated × After	0.079	0.006	0.006	-0.127
	(0.98)	(0.35)	(0.37)	(-0.54)
Running	0.000	-0.000	-0.000	0.002
	(0.19)	(-0.73)	(-0.72)	(1.71)
Running × Treat	0.000	-0.000	-0.000	-0.001
	(0.04)	(-0.37)	(-0.37)	(-0.34)
Running^2 × Treat	-0.000	0.000	0.000	0.000
	(-0.30)	(0.31)	(0.31)	(0.54)
Lender×Time FE	Yes	Yes	Yes	Yes
Vehicle×Time FE	Yes	Yes	Yes	Yes
State×Time FE	Yes	Yes	Yes	Yes
Income×Time FE	Yes	Yes	Yes	
Obs	179,939	179,939	179,939	179,958

NOTE.—This table reports coefficient estimates from Equations 3. The dependent variable is an indicator for whether the borrower has a co-obligor, whether the borrower's employment is verified, whether the borrower's income is verified, and the natural log of the borrower's income. Standard errors are clustered at the running varaible level.
* Significant at the 10% level.
** Significant at the 5% level.
*** Significant at the 1% level.

Table 7: Loan performance: lender-level analysis

Panel A: Loan performance	(1) 12-month default	(2) 24-month default	(3) 36-month default	(4) Interest rate	
Automated Lender × After	0.012**	0.019*	0.026**	0.004**	
	(2.16)	(1.93)	(2.15)	(1.96)	
Lender FE	Y	Y	Y	Y	
Vehicle×Time FE	Y	Y	Y	Y	
State×Time FE	Y	Y	Y	Y	
Income×Time FE	Y	Y	Y	Y	
Credit Score×Time FE	Y	Y	Y	Y	
R^2	0.097	0.154	0.211	0.791	
Obs	4,452,043	4,452,043	3,862,990	4,452,043	

Panel B: Borrower characteristics	(1) log(income)	(2) log(credit score)	(3) Income verification	(4) Employment verification	(5) Co-obligor
Automated Lender × After	-0.007	-0.005	-0.008	0.002	0.004
	(-0.63)	(-1.17)	(-0.99)	(0.29)	(1.34)
Lender FE	Y	Y	Y	Y	Y
Vehicle×Time FE	Y	Y	Y	Y	Y
State×Time FE	Y	Y	Y	Y	Y
Income×Time FE		Y	Y	Y	Y
Credit Score×Time FE	Y		Y	Y	Y
R^2	0.241	0.453	0.316	0.330	0.180
Obs	4,452,060	4,452,043	4,452,043	4,452,043	4,452,043

NOTE.—This table reports coefficient estimates from Equations 6. Outcomes include loan performance, contract terms, and borrower characteristics. The sample is restricted to auto loans originated within two years before and two years after the treatment date. Standard errors are clustered at the lender level.
* Significant at the 10% level.
** Significant at the 5% level.
*** Significant at the 1% level.

Table 8: Difference-in-differences regression: the share of automation

	(3) Automated underwriting	(4) Automated underwriting	(5) Automated underwriting
Automated Lender × After	-0.026**	-0.026**	-0.026**
	(-2.42)	(-2.53)	(-2.53)
Lender FE	Yes	Yes	Yes
Vehicle×Time FE	Yes	Yes	Yes
State×Time FE	Yes	Yes	Yes
Income×Time FE			Yes
Credit Score×Time FE		Yes	Yes
R^2	0.326	0.351	0.351
Obs	4,452,060	4,452,060	4,452,043

NOTE.— This table reports coefficient estimates from Equation 6. The dependent variable is the share of loans originated through automated systems. The sample is restricted to auto loans originated within two years before and two years after the treatment date. Standard errors are clustered at the lender level.

* Significant at the 10% level.

** Significant at the 5% level.

*** Significant at the 1% level.

Table 9: Cross-sectional tests: lender-level analysis

Panel A: Loan performance	(1) 12-month default	(2) 12-month default	(3) 12-month default
Automated Lender × After	0.008*	0.002	0.010**
	(1.85)	(0.96)	(2.10)
Automated Lender × After × Low Income	0.008*		
	(1.86)		
Automated Lender × After × Low CS		0.026**	
		(2.21)	
Automated Lender × After × High LTV			0.005*
			(1.94)
Lender FE	Y	Y	Y
Vehicle×Time FE	Y	Y	Y
State×Time FE	Y	Y	Y
Income×Time FE		Y	Y
Credit Score×Time FE	Y		Y
R^2	0.099	0.093	0.100
Obs	4,451,858	4,451,813	4,452,018

Panel B: Loan pricing	(1) Interest rate	(2) Interest rate	(3) Interest rate
Automated Lender × After	0.001	-0.002***	0.003
	(0.50)	(-3.83)	(1.11)
Automated Lender × After × Low Income	0.005***		
	(4.86)		
Automated Lender × After × Low CS		0.010***	
		(5.35)	
Automated Lender × After × High LTV			0.002
			(0.69)
Lender FE	Y	Y	Y
Vehicle×Time FE	Y	Y	Y
State×Time FE	Y	Y	Y
Income×Time FE		Y	Y
Credit Score×Time FE	Y		Y
R^2	0.793	0.748	0.795
Obs	4,451,858	4,451,813	4,452,018

NOTE.—This table reports coefficient estimates from Equation 7. Panel A uses default within 12 months as the dependent variable, and Panel B uses the interest rate at origination. Column (1) defines Low Income as borrowers with income below the median, Column (2) defines Low CS as borrowers with credit scores below the median, and Column (3) defines High LTV as borrowers with loan-to-value ratios above the median. Standard errors are clustered at the lender level.
* Significant at the 10% level.
** Significant at the 5% level.
*** Significant at the 1% level.

Figure 1: Underwriting process in auto loan market

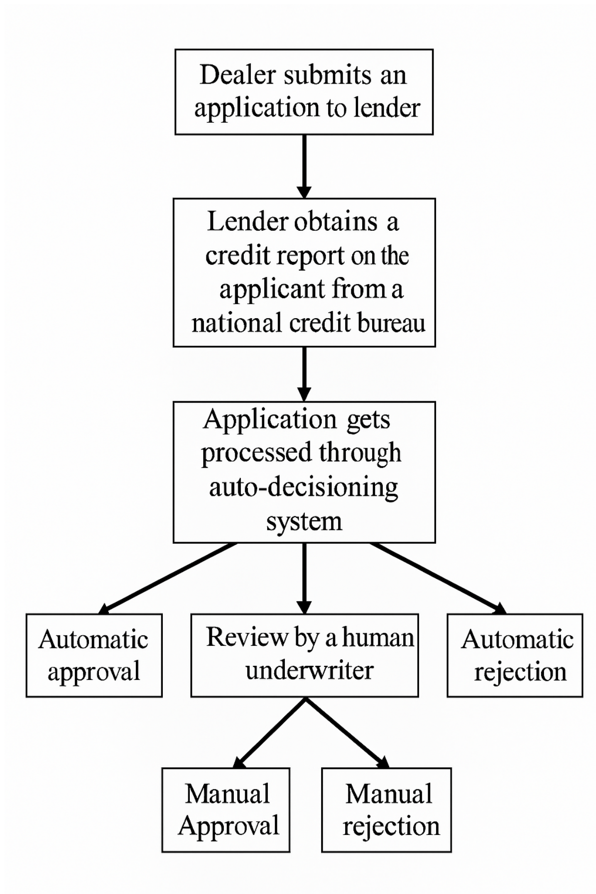

NOTE.——This figure presents underwriting process for lenders with both automated and manaul underwriting in the U.S. auto loan market.

Figure 2: Share of automated and manual underwriting across lenders

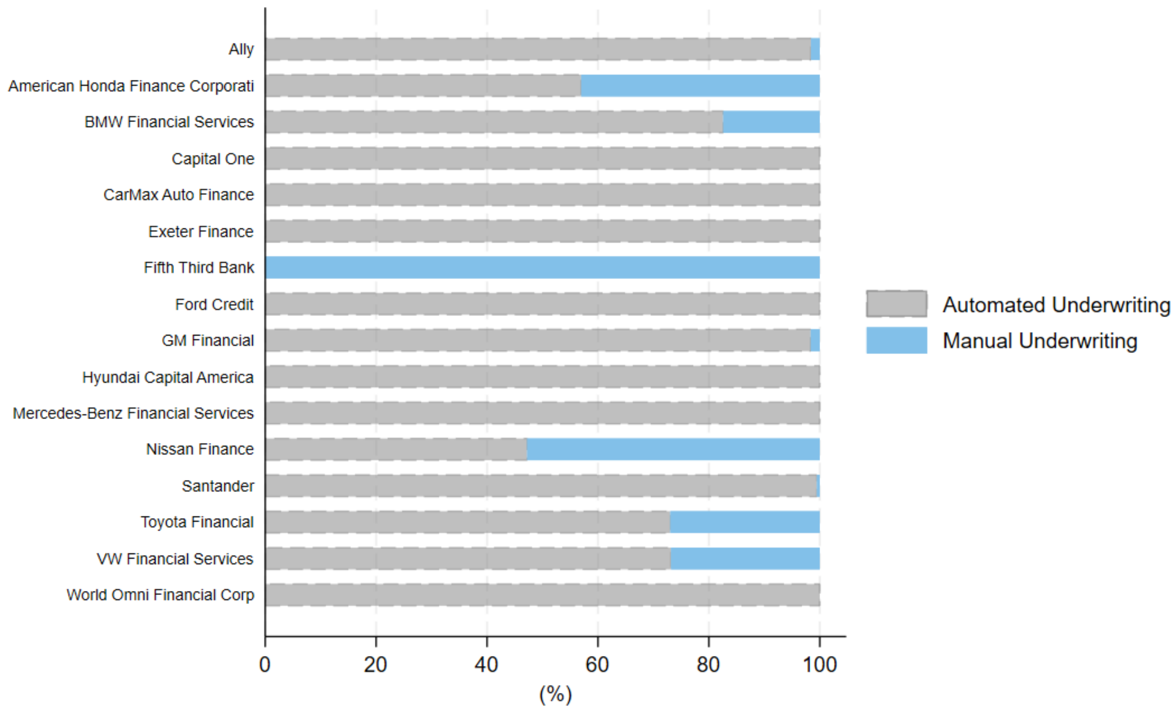

NOTE.——This figure shows the share of automated and manual underwriting across lenders from 2018 to 2022. For each lender, the gray bar represents the percentage of loans originated through automated systems, while the blue bar represents the percentage originated through manual review.

Figure 3: Heterogeneity in automated underwriting

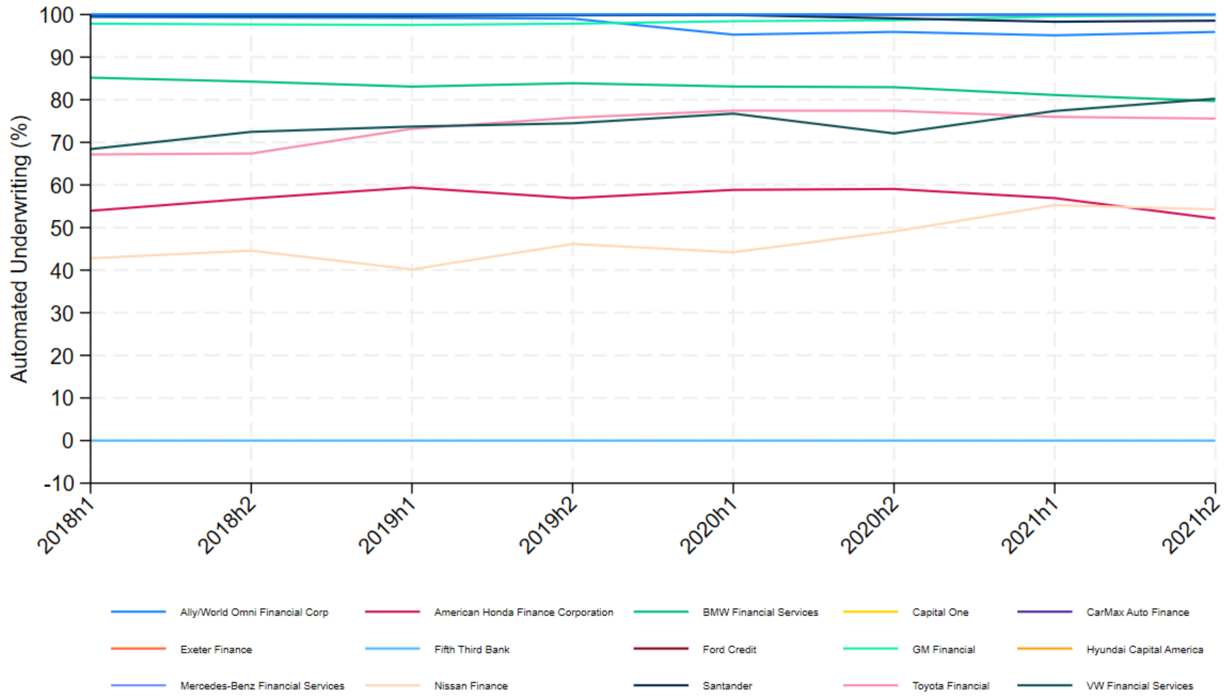

NOTE.——This figure plots the share of loans originated through automated systems from 2018 to 2022, highlighting the stability of automation practices over time.

Figure 4: Discontinuities in the probability of automated underwriting

(a) Nissan

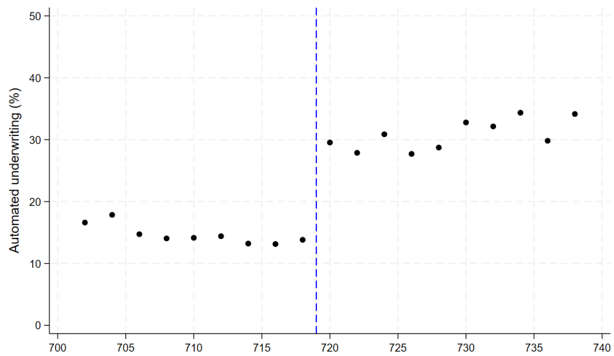

(b) Volkswagen

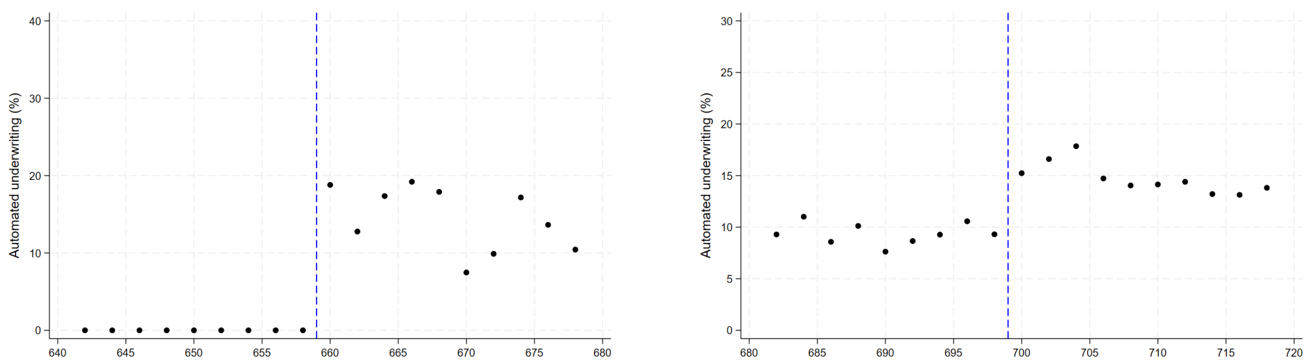

NOTE.—This figure plots the average probability of automated underwriting within 2-point credit score bins for (a) Nissan Finance and (b) Volkswagen Financial Services.

Figure 5: Local continuity of borrower characteristics: Nissan

(a) Income

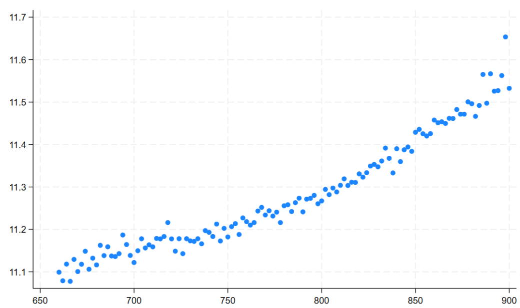

(b) Income verification

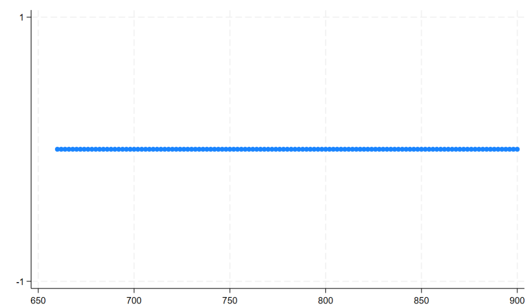

(c) Employment verification

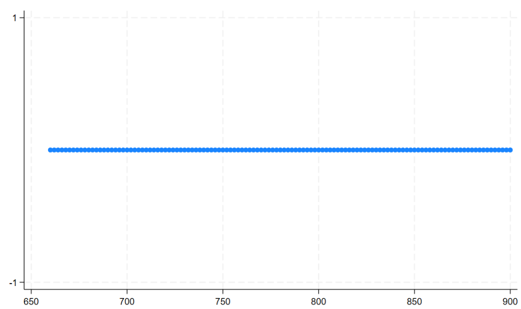

(d) Co-obligor

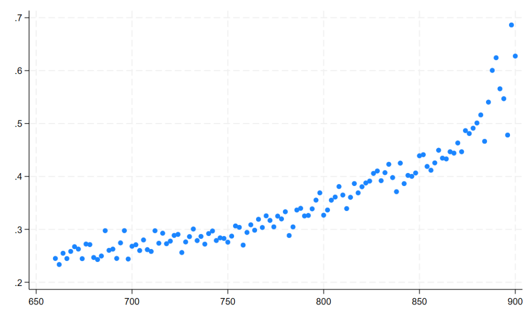

NOTE.——This figure plots raw average of borrower characteristics across credit scores for Nissan. The borrower characteristics are (a) the natural logarithm of income, (b) an indicator for whether the borrower's income is verified, (c) an indicator for whether the borrower's employment is verified, and (d) an indicator for whether the borrower has a co-obligor. Circles represent the average value of each characteristic within 2-point credit score bins.

Figure 6: Local continuity of borrower characteristics: Volkswagen

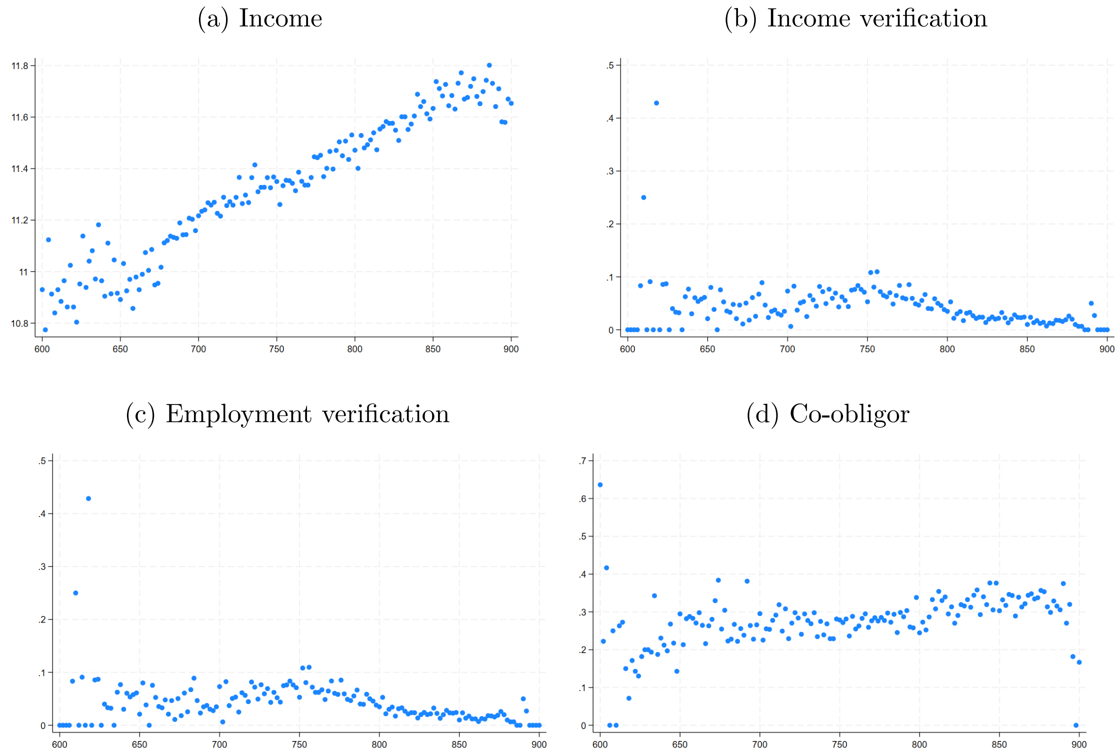

NOTE.——This figure plots raw average of borrower characteristics across credit scores for Volkswagen. The borrower characteristics are (a) the natural logarithm of income, (b) an indicator for whether the borrower's income is verified, (c) an indicator for whether the borrower's employment is verified, and (d) an indicator for whether the borrower has a co-obligor. Circles represent the average value of each characteristic within 2-point credit score bins.

Figure 7: Dynamic Loan Performance: Semiannual Estimates

(a) 12-month default

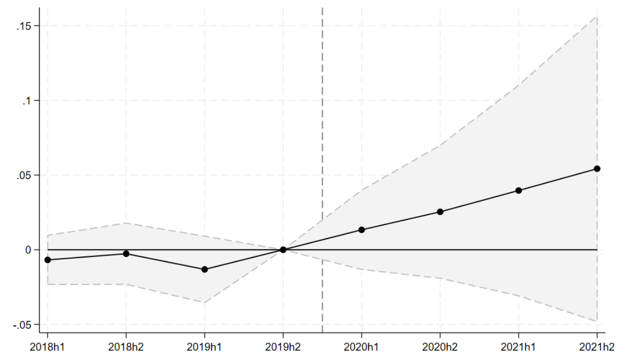

(b) 24-month default

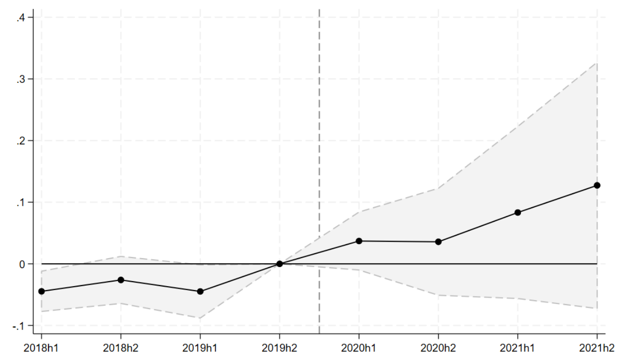

(c) 36-month default

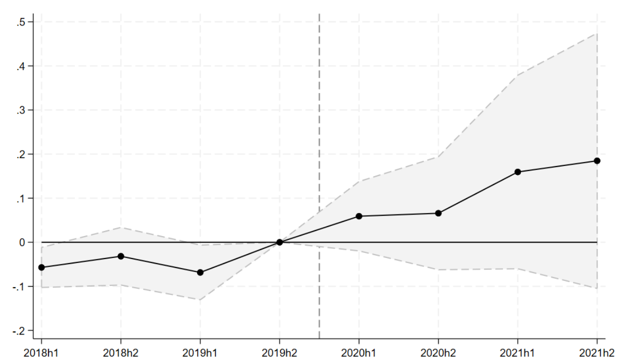

NOTE.——This figure plots coefficient estimates from Equation 4. The dependent variable is either the 12-month default, the 24-month default, or 36-month default. The x-axis corresponds to the number of half-year periods relative to the treatment date. The period $\tau = -1$ is the reference point. The circles correspond to the coefficient estimates, and shaded areas indicate 95 percent confidence intervals. Standard errors are clustered at the running variable level.

Figure 8: Baseline specification: loan terms

(a) Interest Rate

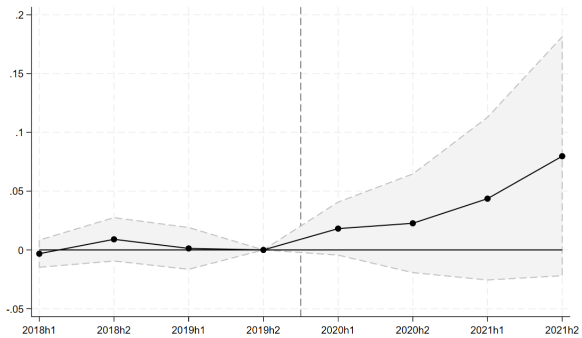

(b) Maturity

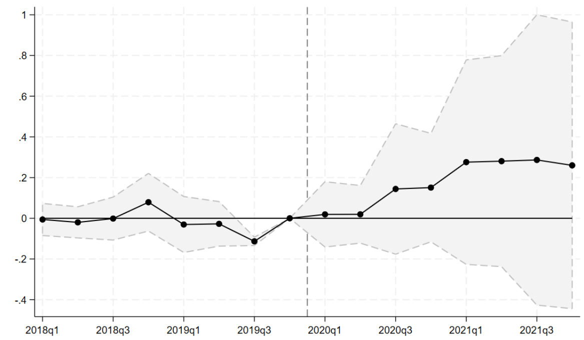

(c) Loan-to-Value

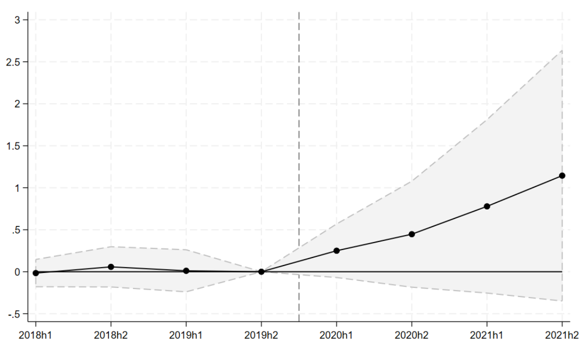

(d) Loan Amount

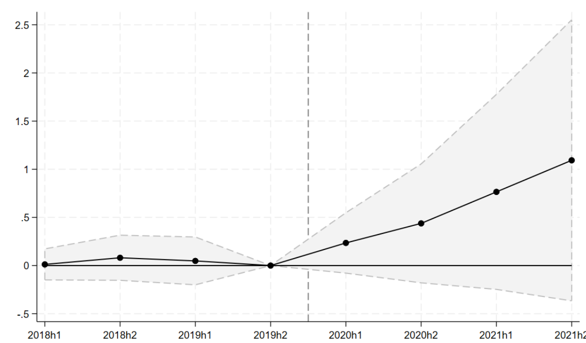

NOTE.——This figure plots coefficient estimates from Equation 4. The loan term is either (a) interest rate, (b) the natural log of the loan maturity, (c) loan-to-value ratio , or (d) the natural log of the loan amount. The x-axis corresponds to the number of half-year periods from the treatment date. The period $\tau = -1$ is the reference point. The circles correspond to the coefficient estimates, and shaded areas indicate 95 percent confidence intervals. Standard errors are clustered at the running variable level.

Figure 9: Baseline specification: borrower characteristics

(a) Income

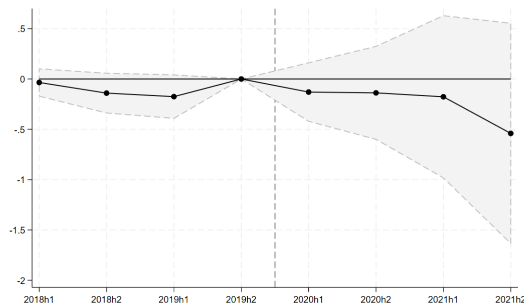

(b) Co-obligor

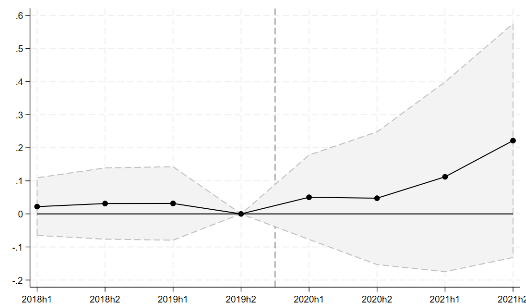

(c) Income verification

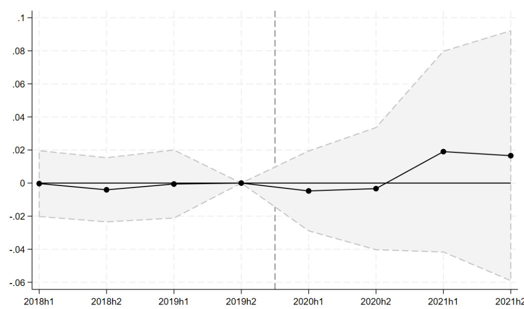

(d) Employment verification

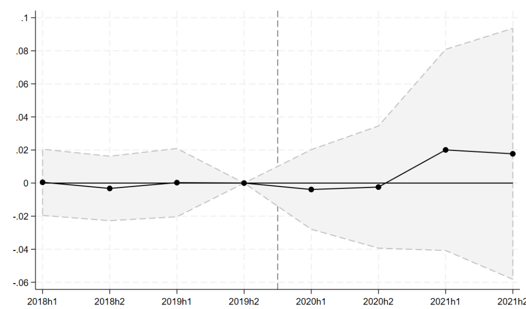

NOTE.——This figure plots coefficient estimates from Equation 4. The borrower characteristic is either (a) the natural log of income (b) the presence of a co-obligor, (c) verification of income, or (d) verification of employment. The x-axis corresponds to the number of half-year periods relative to the treatment date, with $\tau = -1$ serving as the reference point. Circles represent point estimates, and shaded areas indicate 95 percent confidence intervals.

Figure A.1: Discontinuities in probability of automated underwriting

(a) Nissan

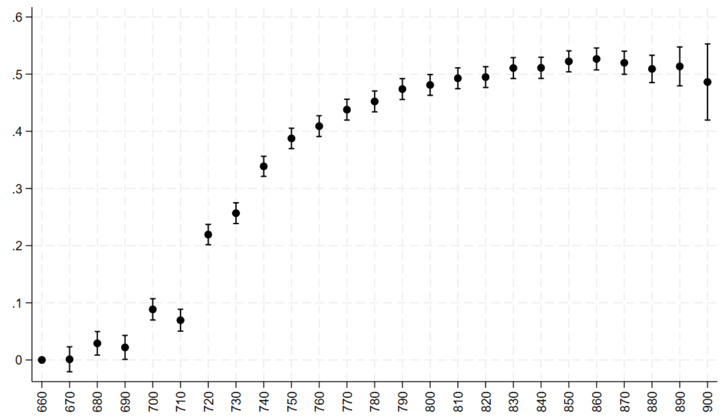

(b) Volkswagen

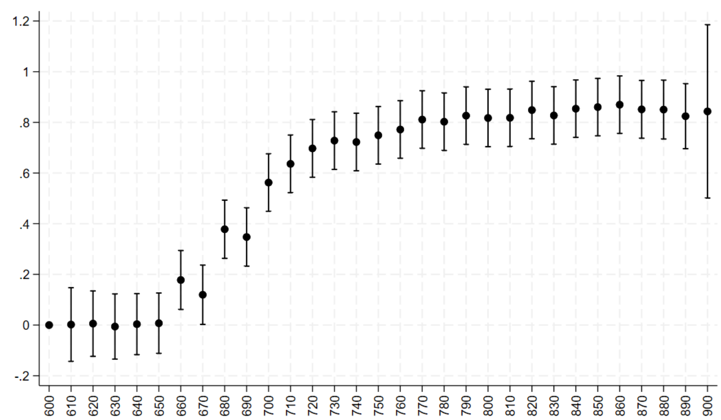

NOTE.——This figure plots the conditional average of automated underwriting across credit score bins for two lenders: (a) Nissan, and (b) Volkswagen. The circles correspond to the coefficient estimates, and the vertical bars indicate 95 percent confidence intervals. Standard errors are clustered at the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

Figure A.2: Dynamic loan performance: semiannual estimates over extended period

(a) 12-month default

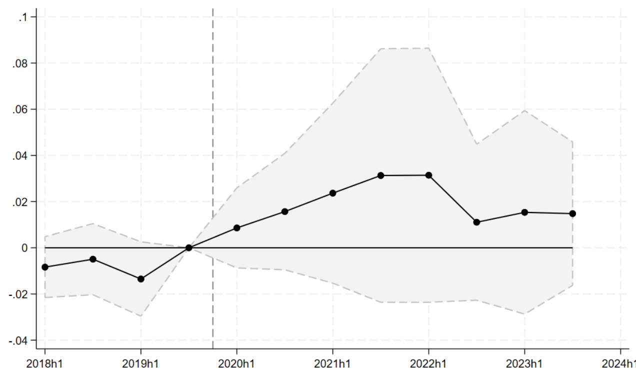

(b) 24-month default

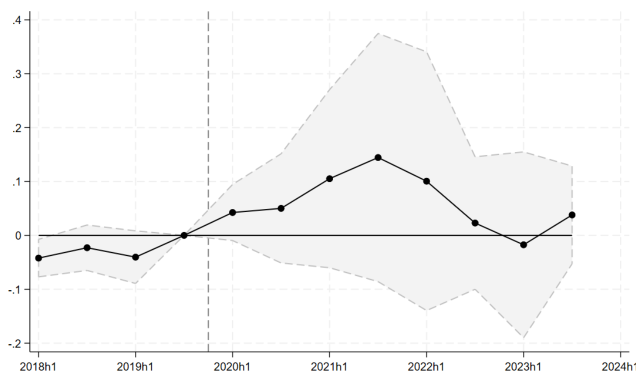

(c) 36-month default

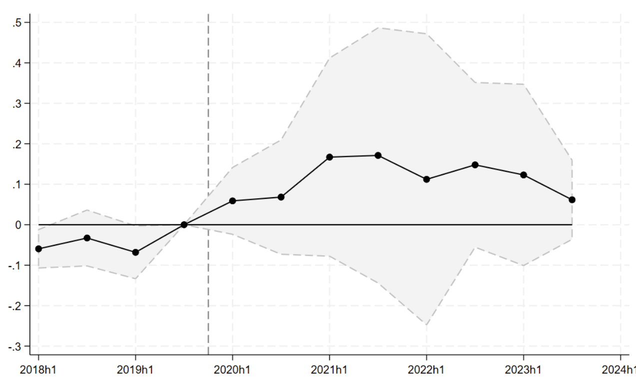

NOTE.——This figure plots coefficient estimates from Equation 4. The dependent variable is either the 12-month default, the 24-month default, or 36-month default. The x-axis corresponds to the number of half-year periods relative to the treatment date. The half-year period $\tau = -1$ is the reference point. The circles correspond to the coefficient estimates, and shaded areas indicate 95 percent confidence intervals. Standard errors are clustered at the running variable level.

Figure A.3: Dynamic loan performance: quarterly estimates

(a) 12-month default

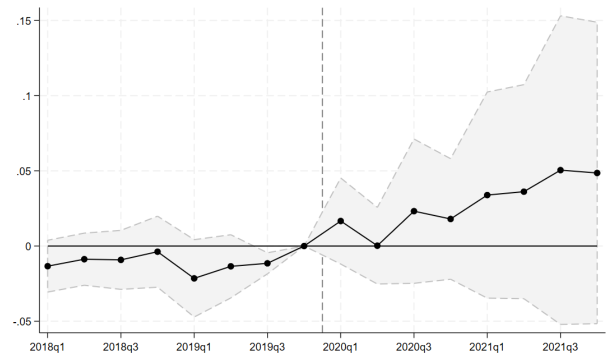

(b) 24-month default

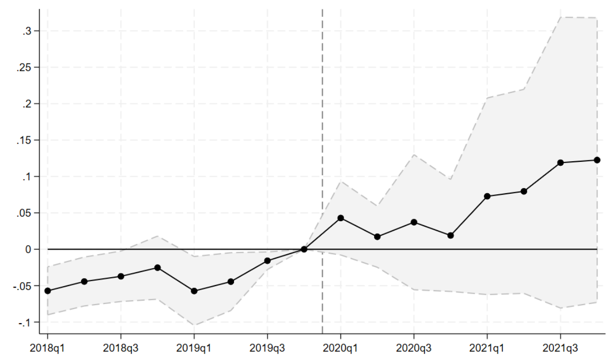

(c) 36-month default

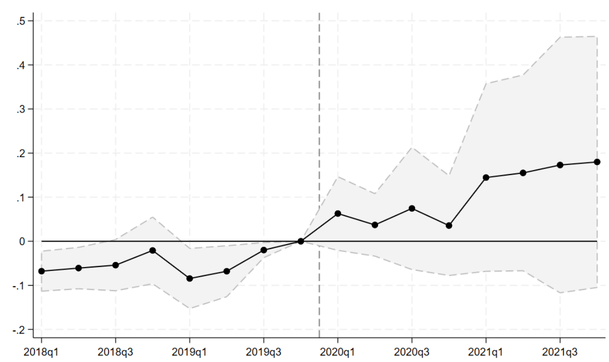

NOTE.——This figure plots coefficient estimates from Equation 4. The dependent variable is either the 12-month default, the 24-month default, or 36-month default. The x-axis corresponds to the number of quarters from the treatment date. The quarter $\tau = -1$ is the reference quarter. The circles correspond to the coefficient estimates, and shaded areas indicate 95 percent confidence intervals. Standard errors are clustered at the running variable level.

Figure A.4: Dynamic loan performance: quarterly estimates over extended period

(a) 12-month default

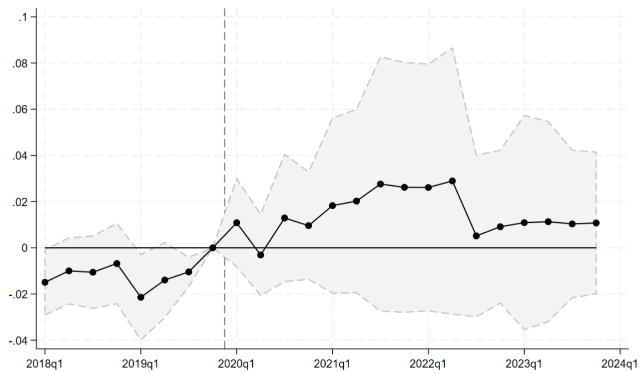

(b) 24-month default

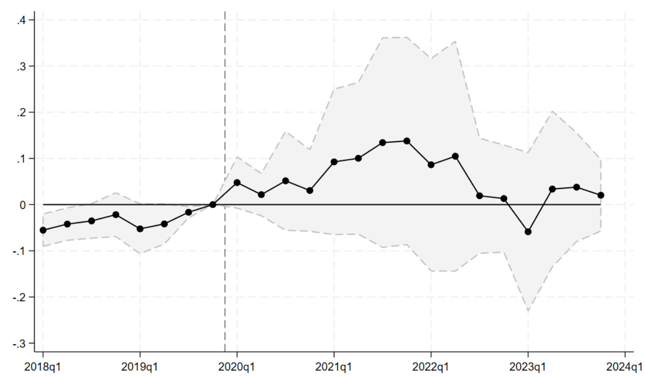

(c) 36-month default

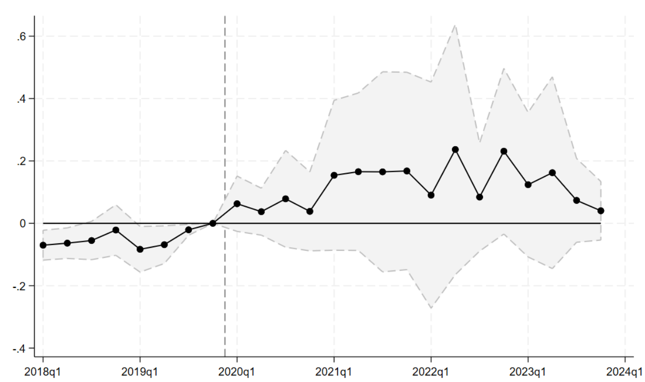

NOTE.——This figure plots coefficient estimates from Equation 4. The dependent variable is either the 12-month default, the 24-month default, or 36-month default. The x-axis corresponds to the number of quarters from the treatment date. The quarter $\tau = -1$ is the reference quarter. The circles correspond to the coefficient estimates, and shaded areas indicate 95 percent confidence intervals. Standard errors are clustered at the running variable level.