

How Does Human–Machine Collaboration Work?

— Evidence From Auto Finance Leasing Transactions*

Yao Lu[†] Hanqing Shi[‡] Yanling Sun[§] Zhishu Yang[¶]

ABSTRACT

This paper studies the dynamic evolution of human behavior and collaborative value—along both efficiency and fairness dimensions—in algorithm-assisted credit approval. Using transaction-level data from a leading Chinese auto finance leasing firm, we analyze loan officers’ override (rescue) decisions for applications initially rejected by a machine-learning credit scoring system. We document three main findings. First, rescued applications shift over time from being broadly distributed across low score ranges to clustering near the algorithmic cutoff, indicating increasingly margin-focused human intervention. Second, human overrides reduce reliance on coarse group-level information, thereby mitigating statistical discrimination, without sacrificing default performance or profitability, with collaborative value shifting from fairness gains early on to risk reduction over time. Third, mechanism analyses point to a career-concern channel: following algorithm adoption, loan officers initially exert greater effort to signal competence, as reflected in richer approval texts and longer processing times; as career concerns attenuate over time, incentives weaken, leading to partial free-riding on the algorithm and reduced human input. Together, these findings highlight the dynamic nature of human–machine collaboration and the importance of organizational incentives in sustaining its value.

JEL Classification: G21, D83, O33

Keywords: Human–Machine Collaboration, Credit Approval, Statistical Discrimination, Career Concerns, In-House Finance

*We are grateful to Lin William Cong, J.Anthony Cookson, Manish Jha, Hogyu Jhang, and Jian Sun for their insightful comments and suggestions. We also thank participants at the 2025 CUHK-RAPS-RCFS Conference, 2025 20th Conference on Asia-Pacific Financial Markets (CAFM), Five-Star Conference in Finance, and Tsinghua University for valuable feedback. All errors are our own.

[†]School of Economics and Management, Tsinghua University, luyao@sem.tsinghua.edu.cn.

[‡]School of Economics and Management, Tsinghua University, shq20@mails.tsinghua.edu.cn.

[§]Graduate School of Management, Politecnico di Milano, yanling.sun@gsom.polimi.it.

[¶]School of Economics and Management, Tsinghua University, yangzhsh@sem.tsinghua.edu.cn.

I. Introduction

The rapid diffusion of artificial intelligence and machine learning has fundamentally reshaped decision-making in high-stakes environments such as credit approval, hiring, medical diagnosis, and judicial risk assessment. A growing literature has begun to systematically compare the performance of human decision-making, algorithmic decision-making, and hybrid human-machine interaction in financial decision-making, examining how outcomes differ when decisions are delegated to humans, algorithms, or their combination (Cao et al., 2024; Lu and Zhang, 2025; Wang et al., 2025; Jansen et al., 2025). However, largely due to data limitations, existing studies adopt a predominantly static perspective, examining human-machine interaction at a single point in time or shortly after algorithm adoption. As a result, little is known about how human behavior, organizational incentives, and learning dynamics evolve as algorithmic systems mature, or whether the role and value of human discretion strengthen, weaken, or qualitatively change over time. Moreover, prior work largely focuses on efficiency or predictive accuracy, despite growing evidence that algorithm adoption can have important distributional implications (Liang et al., 2026). Understanding human-machine collaboration therefore requires a dynamic and joint evaluation of both efficiency and fairness.

Using granular transaction-level data from the auto finance leasing sector over a long horizon (2018–2023), including rich approval texts and detailed processing-time measures, we examine how human decision-makers interact with algorithmic recommendations in a hybrid credit approval process. Our empirical analysis yields three key findings. First, human override behavior evolves over time: early after adoption, rescued applications are broadly distributed across low score ranges, while later rescue activity becomes increasingly concentrated near the algorithmic cutoff, indicating more margin-focused intervention. Second, human overrides reduce reliance on coarse group-level information without increasing default risk or reducing profitability, with collaborative value shifting from fairness gains in the early period to risk reduction over time. Third, mechanism analyses point to a career-concern channel: following algorithm adoption, loan officers initially exert greater effort, reflected in richer approval texts and longer processing times, but as career concerns attenuate, incentives weaken and human input declines. Together, these findings show that human judgment plays a distinct yet evolving role in algorithm-assisted credit approval, highlighting how organizational incentives and principal-agent frictions shape the dynamic value of human-machine collaboration.

The data firm is a subsidiary of one of China’s largest automobile manufacturers, with business operations covering all provinces. It serves as the vendor-financing arm of the parent company’s passenger vehicle division, facilitating automobile purchases through in-house financial leasing solutions. Consequently, the firm operates under dual business objectives: stimulating vehicle sales for the parent company while generating interest income from financial

leasing. This dual mandate distinguishes the firm from traditional bank auto lending, where the sole objective is interest income generation, and creates a fundamental trade-off between expanding the customer base and controlling default risk. Within this institutional context, loan performance is more appropriately characterized by considering both default outcomes and realized profitability.

To improve the efficiency and consistency of credit approval, the firm adopts a FinTech-based credit scoring system developed by a leading technology provider specializing in credit risk analytics. The system applies a gradient boosting algorithm (XGBoost) trained and validated on large-scale auto finance datasets, generating a credit score ranging from 300 to 1000 together with a preliminary approval recommendation—approve if the score exceeds 475 and reject otherwise. Credit officers retain full discretion and make the final decision, using the machine score as a reference. As a result, observed human overrides primarily reflect differences in judgment rather than differences in information. Consistent with this design, approval patterns closely track the score distribution: in the high-score range (above 475), officers follow the algorithm’s recommendation in 92.42% of cases, whereas in the low-score range (475 or below), officers override machine rejections in 59.01% of cases. Importantly, lending decisions are based almost exclusively on standardized, verifiable hard information contained in the loan application, leaving limited scope for credible soft information to affect approval outcomes. This institutional setting provides a clean environment for studying the dynamics of human–machine collaboration in credit decision-making.

We begin by measuring the collaborative value of human–machine interaction along two dimensions: post-approval economic performance and fairness in credit allocation. To capture performance, we go beyond default-based risk measures and jointly evaluate loan repayment outcomes and realized profitability, recognizing that in in-house auto finance leasing, default does not necessarily translate into financial loss. To assess fairness, we examine the extent to which approval decisions substitute coarse group-level attributes for individual credit potential, and quantify this substitution using a coarse-group information reliance index that captures systematic dependence on demographic and socioeconomic group proxies in approval rules.

As a first step in the empirical analysis, before turning to the dynamics of rescued cases, we assess the baseline effects of machine-based credit scoring on borrower outcomes and information use. Exploiting institutional variation from the introduction of the FinTech scoring system across vehicle segments, we show that machine-assisted decisions are associated with significantly lower default risk relative to human-only approval. At the same time, algorithm adoption leads to a systematic increase in reliance on coarse group-level information, indicating a higher degree of statistical discrimination in credit allocation. These results point to a trade-off introduced by algorithmic scoring: while predictive performance improves, approval decisions

place greater weight on group-based statistical signals rather than individual-specific assessments. The findings are robust across alternative specifications, including a complementary difference-in-differences design focusing exclusively on the passenger vehicle segment.

Having established this baseline trade-off, we then examine whether and how human discretion can reshape these outcomes through rescued decisions. Focusing on applications rejected by the machine but approved by credit officers, we show that rescued decisions do not worsen loan performance: default risk and loan profitability remain statistically indistinguishable from loans jointly approved by both human and machine. In contrast, rescued decisions significantly reduce reliance on group-level attributes in approval decisions, as reflected in a lower coarse-group information reliance index after controlling for detailed officer- and agent-level heterogeneity. These results indicate that rescued decisions can partially offset the increased statistical discrimination associated with algorithmic scoring, improving fairness without sacrificing economic performance.

We next study how the human role in rescued decisions evolves by splitting the post-adoption passenger-vehicle sample into two equal-length phases. We find that in the later phase, rescued activity shifts markedly toward the machine’s decision margin: rescued cases become increasingly concentrated near the cutoff, consistent with a progressive convergence of officers’ decisions toward algorithmic recommendations. We then examine whether the value created by rescued decisions changes in tandem with this behavioral convergence. The evidence points to a clear reallocation of collaborative value across dimensions: in the early post-adoption phase, rescued decisions primarily operate through mitigating reliance on coarse group-level proxies, whereas in the later phase they increasingly operate through screening out default risk, with loan profitability remaining largely unaffected.

To interpret these dynamics, we consider two competing mechanisms. Under a trust-learning channel, officers may initially discount algorithmic recommendations but, through repeated exposure and outcome feedback, update beliefs about model competence and gradually move toward calibrated reliance, focusing attention on marginal cases where human judgment is most valuable. Under a career-concern channel in the spirit of Holmström (1999), early adoption heightens replacement risk and reputational uncertainty, inducing officers to exert visible discretionary effort to signal competence; as the system matures and accountability benchmarks stabilize, reputational uncertainty declines, incentives to exert effort weaken, and behavior shifts toward heavier reliance on the algorithm and reduced human input consistent with partial free-riding.

To distinguish between the trust-learning channel and the career-concern channel, we marshal evidence from four complementary dimensions of the decision process. First, we examine officers’ approval-opinion texts using large language model-based methods, including Sentence-

BERT embeddings to quantify semantic similarity and defensiveness. Although these texts do not directly reveal officers' true beliefs, they provide a useful window into how decision rationales are articulated under organizational constraints. We document a systematic dynamic shift in both effort and content: over time, approval opinions become shorter, less information-dense, and more semantically homogeneous, with a markedly higher incidence of vague and template-based expressions. These patterns are especially pronounced in the later post-adoption period and for applications near the algorithmic cutoff, where decisions are most contestable. Rather than becoming more detailed or substantive as officers gain experience with the algorithm, written justifications increasingly shift toward reusable and *ex post* defensible language. This evolution is difficult to reconcile with a calibrated trust-learning view, which would predict richer explanations precisely in marginal cases where human judgment is most valuable. Instead, the observed contraction in textual effort and the rise in defensive phrasing are more consistent with a career-concern mechanism: as reputational uncertainty declines and the algorithm becomes the dominant accountability benchmark, marginal incentives to exert and document independent judgment weaken, leading officers to economize on both decision effort and its written articulation.

We then analyze transaction-level audit time to capture effort provision directly. The time pattern mirrors the textual evidence. Audit time falls substantially as the system matures, and the contraction is not confined to mechanically “easy” cases. In the later post-adoption period, time investment decreases for both machine-approved and machine-rejected applications, with particularly sharp declines around the algorithm's decision boundary. This allocation runs counter to the trust-learning prediction that officers should reallocate scarce attention toward marginal cases where algorithmic uncertainty is greatest, and instead aligns with reduced effort provision and greater reliance on the algorithm as the organizational benchmark for accountability.

Third, we exploit promotion periods—episodes of heightened performance pressure, more salient internal ranking, and tighter links between short-run outcomes and career evaluation—to provide a sharper test of incentive-based explanations. During promotion periods, officers expand rescued activity into lower-score regions and become more willing to deviate from algorithmic rejections, partially reversing the late-stage convergence toward the cutoff observed in regular periods. This incentive-sensitive re-expansion of discretion is difficult to attribute to passive belief updating alone and instead supports a career-concern interpretation in which visible intervention rises when career incentives intensify.

Finally, We further probe the career-concern mechanism using heterogeneity in officers' incentives and uncertainty. Focusing on rescued passenger-vehicle applications, we show that the late-stage, near-cutoff contraction in human effort and articulation is not uniform: audit time

and substantive writing fall more, while justifications become more vague and semantically standardized, precisely for officers and environments where accountability pressures should be stronger. The effects are amplified for officers with steeper career stakes (tenure), for those with more performance-relevant uncertainty (higher volatility in pre-FinTech profit ratio), and under heightened external uncertainty (COVID exposure). In contrast, volatility in pre-FinTech CGIR is comparatively weak and often insignificant, suggesting that what matters most for late-stage defensibility is uncertainty about performance rather than dispersion in fairness-related style.

Building on the above mechanism evidence, we further exploit a key advantage of our data: we observe the complete partition of applications by machine and human decisions, including not only jointly approved and rescued cases, but also applications ultimately rejected. This allows us to extend the analysis beyond rescued loans and separately study upward human overrides within the machine-rejected pool and downward human overrides within the machine-approved pool. Descriptively, both types of deviations from the algorithm become systematically less frequent in the later post-adoption period, consistent with a defensibility logic under the career-concern channel whereby adhering to the algorithm increasingly becomes the safe organizational benchmark.

Using counterfactual outcome estimates to place approved and rejected applications on a common performance scale, we show that upward overrides on average improve selection in the rejected region, yielding lower default risk and higher profitability without worsening fairness, although their contribution shifts over time toward stronger risk screening accompanied by a partial re-emergence of group-based reliance. Downward overrides, in contrast, act as a risk backstop within the approved region by screening out loans with worse risk–return profiles, but their precision deteriorates as the system matures: late-stage vetoes remain conservative in terms of risk reduction yet increasingly forgo profitable lending opportunities, suggesting a shift toward more defensive and coarser human intervention as the algorithm becomes institutionally entrenched.

This paper contributes to the literature in several ways.

First, to the best of our knowledge, this paper is the first to examine the dynamic evolution of human–AI collaboration from a longitudinal perspective. In contrast to the existing literature, which predominantly studies the performance effects of human–algorithm interaction in static settings (Jansen et al., 2025; Lu and Zhang, 2025; Wang et al., 2025), we exploit a long panel of micro-level loan approval data to trace how loan officers’ behavior evolves over time following the introduction of algorithmic decision support. Specifically, we document a systematic shift in human decision-making from early-stage effort intensification and self-verification to later-stage reliance on algorithms and attenuated human effort. These findings extend the classic dynamic incentive framework of career concerns developed by Holmström (1999) to a novel human–AI

collaboration context. Moreover, by embedding algorithm adoption into the loan underwriting process, our results provide new empirical evidence for the literature on agency conflicts and underwriting behavior in credit markets (Udell, 1989; Hertzberg et al., 2010; Heider and Inderst, 2012; Berg, 2015; Agarwal and Ben-David, 2018). Taken together, our evidence suggests that technological adoption does not simply eliminate agency problems; instead, it reshapes their dynamic manifestation over time.

Second, we integrate fairness and efficiency into a unified framework of collaborative value in human–algorithm interaction and provide direct micro-level evidence that the two objectives can be jointly improved in practice. The existing literature on the trade-off between profitability and fairness in financial decision-making has largely focused on the role of algorithms or machine-only decision systems, emphasizing how different types of information inputs affect predictive accuracy and group-level disparities (Lu et al., 2023; Liang et al., 2026; Hu et al., 2025). Much less attention has been paid to how fairness and efficiency jointly evolve within a human–algorithm collaboration framework. While Lu and Zhang (2025) show that human judgment can correct algorithmic errors arising from randomness or structural bias—thereby improving both predictive accuracy and decision fairness—their analysis does not systematically examine distributional consequences across multiple dimensions of fairness. In contrast, exploiting a real-world credit approval setting with algorithmic screening followed by human review, we show that human intervention does not lead to higher default rates, while significantly reducing the reliance of decisions on group-level characteristics. These findings indicate that human discretion can mitigate algorithmic bias without compromising risk-control performance.

Third, this paper offers a more fine-grained and multidimensional characterization of human decision-making in loan approval by combining rich alternative data with large language model–based methods. Specifically, we jointly exploit approval texts and decision-time measures as process-level data to examine loan officers’ judgment logic and effort allocation from the decision process, rather than focusing solely on final approval outcomes. Prior research shows that textual information can reveal decision-makers’ risk perceptions, attention focus, and underlying psychological states (Tetlock et al., 2008; Gentzkow et al., 2019; Hassan et al., 2019), while decision duration serves as an important proxy for attention allocation and deliberation intensity (Hirshleifer et al., 2009). By integrating textual information, time investment, and the algorithm-use context, we uncover dynamic adjustments in loan officers’ information attention and cognitive allocation following the introduction of algorithmic support. These results provide new micro-level evidence on the behavioral mechanisms underlying human–algorithm collaboration in credit decision-making.

Fourth, this paper contributes to the literature on in-house or vendor finance, where firms extend credit to customers through internal financial divisions, often to promote core product

sales. (Brennan et al., 1988) show that firms with market power may offer vendor financing even when bank credit is available, as it allows price discrimination between cash and credit customers. In practice, in-house financing has become increasingly common among firms offering capital-intensive goods and services, with prominent examples including Ford Credit, Apple’s installment programs (e.g., Apple Card and Apple Financing), and IKEA’s financial services. While most subsequent studies in this literature have focused on trade credit (Schwartz, 1974; Petersen and Rajan, 1994; Daripa and Nilsen, 2011; Giannetti et al., 2011; Chod, 2015; Chod et al., 2019), which typically emphasize long-term borrower–lender relationships, our study shifts attention to a setting where such enduring ties are not necessarily present. We enrich this literature by analyzing a firm that pursues a dual business objective—promoting automobile sales for its parent company and generating interest income from leasing. This structure necessitates a strategic trade-off between expanding the customer base and managing default risk.

The remainder of the paper is organized as follows. Section 2 reviews related literature and develops the dynamic mechanisms. Section 3 provides institutional background on the firm and the FinTech rollout. Section 4 describes the data, sample, and variable construction. Section 5 documents the dynamic evolution of rescued decisions and collaborative value. Section 6 conducts channel analysis to distinguish trust learning from career concerns. Section 7 extends the analysis by exploiting the full partition of applications. Section 8 concludes.

II. Literature Review

A. *Human–Machine Collaboration Performance in Loan Approval*

A growing body of literature documents the realized economic consequences and potential benefits of human–intelligent machine collaboration in lending, primarily along two dimensions: efficiency and fairness.

With respect to efficiency, prior studies across a wide range of domains show that machine-assisted decision-making can substantially improve performance in various areas.¹ In the context of the loan approval process, the overall evidence suggests that the adoption of algorithmic decision tools improves loan profitability, reduces default rates, and accelerates processing speed (Fuster et al., 2019; Vallee and Zeng, 2019). However, the relative performance of algorithms vis-à-vis human underwriters, as well as the effectiveness of human–machine interaction, is

¹including hiring decisions (Horton, 2017; Chalfin et al., 2016; Cowgill, 2018; Hoffman et al., 2018; Bai et al., 2022), medical diagnosis (Kleinberg et al., 2015), and judicial bail decisions (Kleinberg et al., 2018). Within the FinTech literature, machine learning and data-mining techniques have been widely applied to stock price and return prediction (Gu et al., 2020; Cao et al., 2024).

not uniform and critically depends on loan characteristics, information environments, and institutional design. For example, Jansen et al. (2025) show that while human and algorithmic underwriting perform similarly in low-risk and low-complexity cases, algorithms substantially outperform humans in high-risk and high-complexity settings. Beyond technological capability alone, recent studies emphasize the role of human responses to algorithmic recommendations in shaping collaborative outcomes: Wang et al. (2025) and Lu and Zhang (2025) demonstrate that human involvement adds value only under specific informational and explanatory conditions that facilitate effective correction of algorithmic errors. Complementarily, Cao et al. (2024) show that although AI dominates in information-rich and transparent environments, humans retain a comparative advantage when decisions rely on tacit knowledge and institutional judgment.

As for fairness, the existing literature does not reach a unanimous conclusion. One strand of research argues that the adoption of intelligent machines in loan approval can promote financial inclusion by reducing biases in individual lending decisions (Guild, 2017). This effect is mainly driven by two mechanisms. First, intelligent algorithms are able to better exploit non-traditional alternative data sources and incorporate “weak signals” that are not captured by conventional credit models. Combined with high-dimensional machine learning techniques, these data improve risk assessment at the individual level and reduce reliance on group-based characteristics, thereby mitigating statistical discrimination (Loufield et al., 2018; Jagtiani and Lemieux, 2019; Li et al., 2024). Second, algorithm-driven lending decisions reduce preference-based discrimination arising from face-to-face human interactions, such as racial and ethnic discrimination (Bartlett et al., 2022), cultural biases (D’acunto et al., 2026), and discrimination against nonlocal borrowers (Chen et al., 2022).

Another strand of the literature focuses on algorithmic bias and cautions that machine learning-based decision-making under uncertainty may itself generate discriminatory or biased outcomes. These concerns highlight systematic disparities in algorithmic performance across social and demographic groups (see reviews in Kleinberg et al. (2018) and Cowgill and Tucker (2020)).² These biases may arise from multiple sources. One explanation is the limited use of observable sensitive attributes (such as gender or ethnicity) in credit decisions, which are often highly correlated with both creditworthiness and other input features, creating inherent fairness constraints (Fuster et al., 2022). Moreover, bias may originate from the training data itself. In

²In other domains, for example, hobbies listed on a CV may differ by gender, leading algorithms to inadvertently replicate gender biases in hiring decisions (Huang et al., 2024). In online advertising, Google’s machine learning algorithms have been shown to deliver a higher share of “arrest”-related ads for searches of Black names relative to White names (Sweeney, 2013). In healthcare, patients assigned the same risk score by medical algorithms may exhibit significantly different true health risks across racial groups (Obermeyer et al., 2019). In criminal justice, recidivism prediction algorithms generate false-positive rates for Black defendants that are nearly twice those for White defendants (Angwin et al., 2022). Facial recognition technologies also display substantial accuracy disparities across demographic groups (Klare et al., 2012).

credit risk prediction, sample selection remains a relatively underexplored issue: most existing studies rely on samples of approved “good” loans, which can introduce systematic bias when trained models are applied to the full pool of applicants (Cowgill, 2019).

In terms of the trade-off between fairness and efficiency, a growing literature has attempted to integrate the two objectives into a unified analytical framework. For example, Lu et al. (2023) show that different types of alternative data generate heterogeneous trade-offs between profitability and financial inclusion in credit risk assessment. Liang et al. (2026) demonstrate that the fairness–accuracy trade-off is not intrinsic to algorithms but is endogenously driven by group-skewed inputs. Hu et al. (2025) find that machine learning can partially mitigate inherited human biases, with fairness gains varying by the type of bias and borrower history. Despite these advances, relatively little is known about how fairness and efficiency can be jointly improved in settings where humans and algorithms interact in real time. In particular, existing studies largely abstract away from the dynamic nature of human–machine collaboration and offer limited guidance on operational mechanisms that practitioners can deploy to balance fairness and efficiency in practice.

Despite these advances, important questions remain regarding how fairness and efficiency can be jointly improved in settings where humans and algorithms interact in real time. In particular, the existing literature predominantly adopts a static perspective and abstracts from the dynamic evolution of human–machine interaction. To address this gap, our study adopts a dynamic perspective and takes *collaborative value*—the joint improvement in efficiency and fairness generated through human–machine interaction—as the central organizing concept.

B. Dynamic Channels of Human Role in Human–Algorithm Interaction

Prior studies suggest several mechanisms through which the human role in human–algorithm interaction may evolve dynamically over time. We briefly discuss these potential channels and their theoretical foundations below.

B.1. Trust Learning Channel

A central channel emphasized in the literature is trust learning. At the outset of collaboration, humans often approach algorithmic recommendations with skepticism, discounting algorithmic accuracy or privileging their own judgment. Through repeated interaction, decision-makers observe algorithmic performance, receive outcome feedback, and gradually update their beliefs about algorithmic competence. As experience accumulates and information becomes more transparent, initial distrust may give way to calibrated reliance, leading trust in algorithms—and its behavioral implications—to emerge endogenously over time as a key determinant of the

effectiveness of human–algorithm collaboration.

A large body of research on algorithm aversion shows that, even though algorithms outperform humans in a wide range of tasks (Fügener et al., 2021; Krakowski et al., 2023); see also Grove et al. (2000), humans often resist following algorithmic recommendations at early stages of collaboration. Such resistance is commonly attributed to distrust in machines (Jacovi et al., 2021), limited interpretability or lack of reasoning behind algorithmic outputs (Lebovitz et al., 2021), and disappointment when early algorithmic performance fails to meet expectations (Goodyear et al., 2019), all of which can lead decision-makers to disregard algorithmic advice (Dietvorst et al., 2015). Importantly, this literature also emphasizes that algorithm aversion is not a fixed attitude. As users accumulate experience and gain familiarity with how algorithms operate and generate predictions (Luo et al., 2021; Tong et al., 2021), and as they learn to flexibly adjust or correct algorithmic outputs through repeated interactions (Dietvorst et al., 2018), resistance to algorithms tends to attenuate. In parallel, greater algorithmic transparency and the disclosure of decision rationales have been shown to further facilitate trust formation and reduce aversion (Buçinca et al., 2021; Castelo et al., 2019; Kumar and Sharma, 2022; Lu and Zhang, 2025).

Taken together, from a dynamic perspective, human attitudes toward algorithms evolve endogenously over time as a function of usage experience, information transparency, and interaction design. This evolving trust structure, rather than a one-time acceptance decision, plays a central role in shaping the realized performance of human–algorithm collaboration (Fügener et al., 2021; Lu et al., 2023; Wang et al., 2025).

B.2. Career-Concern Channel

Beyond cognitive learning and belief updating, the dynamic evolution of human–algorithm collaboration may also be driven by career incentives and reputational concerns faced by human decision-makers. This mechanism can be understood through the lens of the career-concerns framework developed by (Holmström, 1999). When algorithms are first introduced and relative human versus algorithmic advantages are still uncertain, decision-makers may perceive a heightened risk of replacement and reputational volatility. In response, they tend to exert greater effort by actively intervening in decisions or insisting on human judgment in order to signal competence and “prove their value.” As algorithmic performance stabilizes over time and organizational beliefs about human ability become more fixed, reputational uncertainty declines, weakening marginal incentives to exert effort. This dynamic may manifest as greater reliance on algorithms, attenuated accountability, or even a decline in human input consistent with free-riding behavior.

Existing evidence lends support to this incentive–behavior pathway. First, resistance to

algorithmic recommendations is not driven solely by preferences or psychological aversion, but is closely linked to opportunity costs and economic incentives. When following algorithmic advice entails foregone gains, or when peer behavior conflicts with algorithmic recommendations, individuals are more likely to deviate from algorithmic decisions (Liu et al., 2026). Second, concerns about being replaced by algorithms constitute an important source of career incentives. As artificial intelligence increasingly threatens human labor across multiple domains, decision-makers may strategically emphasize autonomous judgment and highlight human irreplaceability to mitigate perceived substitution risks (Brynjolfsson and Mitchell, 2017; Autor and Dorn, 2013; Siau and Wang, 2018). This logic resonates with the broader “man versus machine” narrative in the literature on job displacement, adaptation, and task reallocation induced by technological change (Acemoglu et al., 2022).

At the same time, the literature highlights a countervailing risk. When tasks are persistently or excessively delegated to algorithms, human decision-makers may gradually reduce their own effort, leading to weakened accountability and diminished complementarity between humans and machines (Fügener et al., 2021). Taken together, these findings point to a dynamic incentive mechanism centered on career concerns: depending on the stage of adoption and the institutional environment, humans may initially intensify effort and resist algorithms in response to substitution threats and high opportunity costs, but subsequently scale back effort and become overly reliant on algorithms as perceived risks decline and convenience increases.

B.3. Task Division Channel

As human–algorithm collaboration unfolds over time, the division of labor between humans and algorithms in information processing and judgment is not static. Instead, decision-makers may gradually delegate standardized and rule-based evaluations to algorithms, while reallocating their limited attention toward dimensions in which algorithms are relatively less effective or less transparent, such as anomaly detection, interpretation of soft information, and value-based trade-offs. This evolving cognitive division of labor reshapes not only when and how humans intervene in decisions, but also the overall performance of human–algorithm collaboration.

A large body of research documents systematic differences in comparative advantage between humans and algorithms across tasks and information environments. Algorithms tend to excel in structured and information-rich settings (Goh et al., 2021), whereas humans perform better in low-frequency events, boundary cases, and highly uncertain environments (Sawyer, 1966; Guo and Wang, 2015), and display greater adaptability in innovative tasks (Lou and Wu, 2021). Related studies further show that when decision tasks are perceived as highly complex, humans are more likely to defer to algorithmic recommendations rather than rely on their own judgment (von Walter et al., 2022), implying that patterns of human–algorithm interaction adjust

systematically with information complexity. More recent work emphasizes that collaborative value does not arise from mechanically following algorithmic advice, but from leveraging humans' general intelligence and cognitive diversity (Te'eni et al., 2026; Wang et al., 2023; Zhang et al., 2021), as well as integrating private or internal information that algorithms cannot easily access (Choudhury et al., 2020; Ibrahim et al., 2021; Sun et al., 2022), thereby forming effective complementarities with algorithmic computational strength (Davenport et al., 2020; Cao et al., 2024). From this perspective, the dynamic nature of human–algorithm interaction can be understood as a continuous process of task and cognitive reallocation: across varying information types, transparency levels, and task complexities, the relative roles, dominance, and intervention strategies of humans and algorithms evolve, ultimately determining whether collaboration generates sustained incremental performance gains.

The institutional setting examined in this study, however, differs fundamentally from the task-division mechanism emphasized above. In our context, algorithms do not directly take over any approval tasks; instead, they function as decision-support tools that provide recommendations to loan officers, while the final lending decision is always made independently by human evaluators. In other words, the introduction of algorithms does not alter the formal allocation of decision authority, but is embedded within the existing human decision process, primarily affecting how loan officers interpret, weigh, and respond to algorithmic information. Given this institutional feature, our analysis does not rely on the task-division channel, but instead focuses on the dynamic mechanisms characterized by the first two channels.

III. Institutional Background

A. The Auto Finance Leasing Market in China

China has maintained its position as the world's largest automobile market for 16 consecutive years in terms of both production and sales volume. In 2024, domestic automobile sales in China reached 25.577 million units, which indicates substantial consumption potential.³ Accompanying the rapid expansion of the auto finance market, the penetration rate of auto finance—defined as the proportion of vehicles purchased using auto finance to total automobile sales—has shown a significant upward trend in recent years. Specifically, the penetration rate of auto finance in China's new car market reached 56% in 2023, representing an increase of approximately 20 percentage points compared to five years earlier.⁴

Auto financial leasing constitutes a critical segment of the broader auto finance market,

³Source: [China Auto Finance Company Industry Development Report](#), published by China Banking Association.

⁴Source: [2024 China Auto Finance Report](#), jointly published by WeBank and Roland Berger.

typically involving three parties: consumers, financial leasing firms, and automobile dealers. The arrangement operates under a rent-to-own model, wherein the leasing company acquires the vehicle from the dealer and retains legal ownership throughout the lease period. The consumer gains usage rights by making regular rental payments, and upon completion of the lease term, ownership is generally transferred to the consumer. This structure lowers the upfront financial burden for consumers and facilitates access to vehicle ownership, particularly for individuals with limited liquidity or credit access.

As non-deposit financial institutions, auto financial leasing companies in China are regulated by the State Financial Regulatory Administration. The regulatory framework includes specific provisions on loan issuance, related-party transactions, and reserve requirements. At the regional level, provincial governments are tasked with developing supportive policies, overseeing risk management, and guiding the industry's growth. Provincial financial regulatory departments are responsible for the day-to-day supervision and administration of financial leasing companies within their jurisdictions.

B. The Business of the Data Firm

The financial leasing company examined in this study is a wholly state-owned subsidiary of a leading automobile manufacturer. It functions as the vendor financing arm of the parent company's passenger vehicle division, providing in-house financial solutions to facilitate automobile purchases. As one of the earliest entrants in China's auto finance market, the company has established a substantial market share, with operations spanning multiple cities. Its key financial indicators remain stable and robust. The firm's funding sources include shareholder capital injections, bank loans, intra-sector leasing arrangements, and asset securitization.

The firm's core financial leasing business is classified into commercial vehicles and passenger vehicles according to the intended purpose of vehicle use. A sample of vehicles in the dataset is shown in Figure A1. Commercial vehicles are designed primarily for the transportation of goods and include buses, trucks, and semi-tractors. Their users are typically individual operators engaged in logistics and transportation services, such as freight and delivery drivers. Passenger vehicles are mainly used for carrying individuals and their personal belongings or small items, encompassing sedans, sport utility vehicles (SUVs), and similar models. These vehicles are primarily leased by household users or natural persons. Across both segments, the firm's clientele consists largely of low- to middle-income individuals who are often excluded from traditional bank credit markets due to insufficient collateral or a lack of formal credit history. As a result, they rely on non-bank financial institutions, such as leasing firms, to obtain credit. This customer segment generally exhibits limited observable creditworthiness and constitutes the long-tail of the consumer credit market.

Unlike traditional bank auto loans that primarily aim to generate interest income, the firm operates under a dual objective: first, to promote automobile sales for its parent company; and second, to earn interest income through auto financial leasing. This dual mandate creates a strategic trade-off between expanding the customer base and managing default risk, reflecting the firm's dual role as both a sales facilitator and a financial intermediary. Consistent with this institutional design, credit officers responsible for loan approvals face performance incentives that extend beyond delinquency control. Their evaluations also reflect the sales performance of the auto business division to which the credit review department is affiliated, thereby aligning credit assessment with broader commercial objectives.

In addition to its distinct business objectives, several institutional features further differentiate the firm's leasing model from traditional bank auto lending: (1) loan maturity—bank auto loans are typically short-term (1–3 years), whereas leasing contracts often extend to 4–5 years; (2) product design—while both banks and leasing companies offer standardized products, the latter provides a more diverse product portfolio and adapts offerings more rapidly in response to market dynamics; (3) regulation—banks, as licensed financial institutions, are subject to stricter regulatory oversight than non-bank leasing firms; (4) collection practices—leasing firms employ more aggressive collection strategies, including in-person recovery efforts, while banks primarily rely on phone-based collection due to their larger loan volumes; and (5) credit evaluation—banks have access to centralized credit bureaus and rely on more automated approval systems, whereas leasing firms often work with limited documentation and use manual assessments.

To control credit risk, the firm adopts a tiered default handling system that escalates with severity and frequency. As shown in Table A1, initial defaults trigger basic actions like phone and app reminders. Repeated or serious delinquencies lead to stronger measures, including credit reporting, vehicle locking, home visits, repossession, and legal action.

C. Credit Approval Process before Machine Adoption

Prior to the adoption of FinTech tools, the credit approval process proceeds as follows. First, a customer selects a vehicle at prevailing market prices from one of the parent company's authorized 4S dealerships and submits a loan application to the firm. Based on the firm's internal credit review guidelines (see Figure A2), a front-line credit officer conducts a preliminary review, which primarily focuses on verifying the authenticity and completeness of the information provided by the applicant. A senior credit officer then performs the final review and makes the ultimate loan approval decision. Because the preliminary review is largely procedural in nature, our analysis focuses on senior credit officers, who are responsible for the final lending decisions. Throughout the sample period, regardless of whether FinTech tools are introduced, each credit application received by the firm is randomly assigned to an individual credit officer.

The final approval decision is based exclusively on the information contained in the submitted application materials. These materials include borrower characteristics, vehicle information, financing terms, supporting application documents, credit information (including court enforcement records and credit bureau reports), and the intended use of the vehicle. Importantly, the firm does not provide explicit approval or rejection rules, nor does it employ shadow scorecards to constrain or guide human discretion. As a result, credit officers rely largely on their experience and judgment when evaluating applications.

In marginal cases, dealerships can provide additional information to strengthen a borrower's application. Specifically, dealerships can include supplementary information in the application's comment section that goes beyond standard credit attributes, with the intention of increasing the likelihood of approval. In our sample, approximately 21.3% of loan applications contain such comments. These comments typically describe recent changes in the borrower's creditworthiness, such as updates to employment status or corrections to inaccuracies in the credit report that have not yet been reflected in credit bureau records. Consistent with the findings of Adelino et al. (2019) and Jansen et al. (2025), we find in unreported results that the presence of comments is not significantly correlated with loan approval outcomes. This suggests that credit officers place little weight on such supplementary information when making lending decisions. The absence of credible soft information allows us to study credit decisions in a clean identification environment, in which observed outcomes are driven primarily by hard information contained in the loan application.

The loan product is selected by the customer prior to submitting the financial leasing application. This selection does not influence the information collected or the approval decision made by credit auditors. To accommodate different customer preferences and enhance process standardization, the company offers a variety of pre-defined loan products. These products vary in terms of down payment ratios, interest rates, loan terms, and payment schedules, allowing for both applicant flexibility and streamlined supervision.

After loan origination, the firm collects monthly payments from borrowers and retains ownership rights to the vehicle until the loan is fully repaid. To control credit risk, the firm adopts a tiered default handling system that escalates with the severity and frequency of delinquency. As shown in Table A1, initial defaults trigger basic actions such as phone calls and app-based payment reminders. Repeated or more serious delinquencies lead to progressively stronger measures, including credit reporting, vehicle locking, home visits, repossession, and legal action. Specifically, when a loan reaches the stage classified as strong measures, the firm attempts to repossess the vehicle and sell it at a public auction. If the auction proceeds, net of recovery costs, exceed the outstanding loan balance, the remaining amount is returned to the borrower. If the proceeds are insufficient to cover the outstanding balance, the lender may file a lawsuit against

the borrower for the deficiency and, upon obtaining a favorable judgment, attempt to recover additional funds through measures such as wage garnishment.

D. Quasi-experiments on Machine Adoption

Since 2020, the company has integrated a FinTech-based credit scoring system developed by an external technology provider into its lending decision-making process. The system generates structured, data-driven score recommendations to support credit officers' assessments, effectively replacing the previous approval framework that relied entirely on officer discretion and subjective judgment.

The implementation of the FinTech scoring system occurs in two phases. Credit scores become available for commercial vehicle applicants starting on February 28, 2020, and for passenger vehicle applicants starting on June 20, 2020. Importantly, the timing of adoption is exogenous to local market conditions, as it is determined centrally by the company's headquarters without regard to regional borrower characteristics. This staggered rollout provides plausibly exogenous policy variation and constitutes a suitable quasi-experimental setting for applying a staggered difference-in-differences design in this study.

The technology provider applies machine learning algorithms, including logistic regression and XGBoost, to model and predict borrowers' repayment ability. The training sample uses historical loan application data and realized default outcomes from the financial leasing company. According to the provider, the scoring model exhibits strong predictive performance and robustness, having been validated in practice across multiple datasets from the auto finance sector.

With respect to the information set, the credit score is generated based on a high-dimensional set of customer data (see Table A2), including identity verification, application information, judicial records, lending history, and other behavioral indicators. These hard information variables fully encompass the information used by credit officers prior to the adoption of FinTech tools and additionally incorporate a richer set of variables. Importantly, the same information is also accessible to loan auditors in the post-adoption period. Specifically, together with the FinTech score, credit officers receive a detailed multidimensional evaluation report that contains both raw input data and algorithm-generated risk alerts. As a result, following the introduction of FinTech, credit officers base their decisions on a strictly expanded set of hard information. However, this information set does not differ between human decision-makers and the algorithm. Our empirical analysis therefore compares lending decisions with and without algorithmic assistance under an identical information environment, allowing us to isolate the effect of machine adoption from differences in available information.

This credit score ranges from 300 to 1000, with higher values indicating lower credit risk and

a lower probability of default. The score serves as an important reference for credit auditors when evaluating loan applications. Based on the score, applicants are classified into two categories: *approve* (scores between 476 and 1000) and *reject* (scores between 300 and 475). Throughout the sample period, the approval threshold remains fixed at 475 points. Figure A3 compares the decision-making principles of human credit officers and the machine-based scoring system.

In addition to the numeric score, the technology provider supplies a detailed multidimensional evaluation report for each applicant, which is accessible to credit auditors. Figure 1 presents the first page of the report, which summarizes the credit score and the overall recommendation. The subsequent pages provide highly detailed analyses of borrower characteristics across multiple dimensions. Each report typically consists of 5 to 7 pages and contains approximately 3,000 to 5,000 words.

Importantly, while credit officers are encouraged to incorporate the FinTech credit score and the accompanying evaluation report into their decision-making process, they are not required to strictly follow the system’s recommendation. Nevertheless, a written justification is required for every loan decision. In particular, if an officer rejects an application with a high credit score or approves one with a particularly low score, the officer must provide a detailed explanation as part of the approval documentation.

Therefore, after the introduction of the machine scoring system, the credit approval process remains largely unchanged and continues to follow a structured sequence that closely resembles the pre-FinTech process. The only substantive difference is that, prior to the senior credit officer’s final review, the FinTech system generates a credit score report together with an approval recommendation. The senior credit officer then conducts the final review, taking into account the FinTech score and its recommendation, and makes the ultimate loan decision. The detailed approval process is illustrated in Figure 2.

IV. Data and Variable

A. Data and Sample

Our proprietary transaction-level dataset covers the full universe of loan applications processed by the auto finance leasing firm described in Section III, spanning the period from January 2018 to December 2023. Relative to prior studies on human–machine interaction in credit approval (Jansen et al., 2025; Lu and Zhang, 2025; Wang et al., 2025; Hu et al., 2025), our dataset offers three key advantages. First, the sample period is unusually long, particularly in the post-FinTech adoption phase, which allows us to examine the dynamic evolution of human–machine collaboration over time rather than static treatment effects. Second, the dataset

includes both approved and rejected passenger vehicle loan applications, enabling us to analyze the respective roles of human discretion and machine recommendations across different decision outcomes. Third, for every application, we observe detailed measures of the time spent by credit officers on approval decisions, as well as rich textual information documenting the stated approval or rejection reasons.

The dataset comprises eight main categories of information. (1) *Loan-level information*, including approval and termination dates, contract region, approved loan amount, down payment ratio, monthly interest rate, and default-related outcomes; (2) *Vehicle-level information*, such as vehicle brand, assessed value, and whether the vehicle is new or used; (3) *Borrower characteristics*, including age, gender, marital status, and place of origin; (4) *Loan product attributes*, such as the product introduction date; (5) *Credit auditor-level information*, covering demographic characteristics (e.g., age, gender, education, and prior work experience) as well as detailed daily performance metrics, including approval efficiency; (6) *Machine-generated credit score data*, produced by the firm's FinTech scoring system; (7) *Decision-processing measures*, including the time spent by credit officers on each application; (8) *Textual approval information*, consisting of written explanations documenting approval or rejection decisions. The first four categories of data—loan-level information, vehicle-level information, borrower characteristics, and loan product attributes—are available for both commercial and passenger vehicle loans. In contrast, the remaining four categories are available exclusively for passenger vehicle applications.

During the data cleaning process, we apply a series of sample selection criteria to ensure the consistency of the analysis. Specifically, we exclude the following observations: (1) duplicate applications submitted multiple times due to system-related errors; (2) applicants who are under the age of 18 at the time of application or over the age of 70 at the end of the loan term; (3) applications involving the purchase of multiple vehicles in a single transaction; (4) applicants identified as corporate entities—we restrict the sample to individual (natural person) applicants only; (5) loan contracts with a duration of fewer than 30 days; (6) commercial vehicle applications dated February 28–29, 2020, and passenger vehicle applications dated June 20–30, 2020, which fall within the initial rollout windows of the FinTech scoring system to ensure consistency in the following monthly-level analysis.

One potential concern relates to the accuracy and reliability of applicant-provided information. We argue that the information submitted to the financial leasing company is generally credible, for three main reasons. First, all applicants are required to submit an official personal credit report issued by the Credit Investigation Center of the People's Bank of China. This report includes verified information on the applicant's identity, residential address, employment, and historical credit behavior. As such, it covers a substantial portion of the data required by the financial leasing company and serves as an authoritative verification source. Second, as described

earlier, the initial credit review process includes dedicated steps for authenticity verification. These procedures involve direct contact with the applicant’s employer to validate declared income and employment status, as well as the use of third-party data sources to cross-check submitted information. Third, we empirically examine whether applicants engage in strategic misreporting by comparing submitted information across multiple applications. Specifically, we analyze whether applicants whose previous loan applications are rejected systematically alter key personal details in subsequent submissions. We find no evidence of systematic manipulation, suggesting that applicants do not, on average, engage in strategic information fabrication following a rejection.

B. Variable Construction

We focus on two primary outcome variables central to our analysis: post-approval performance and the degree of statistical discrimination, measured by a coarse-group information reliance index.

Post-approval performance. We evaluate post-approval loan performance using not only traditional measures of credit risk, such as default rates, but also loan profitability. This distinction is particularly important in the context of in-house auto finance leasing. In this setting, loan default does not necessarily imply a financial loss. First, from a firm-wide perspective, the company generates revenue from the vehicle manufacturing and sales process in addition to lending income. Second, even in the event of default, lenders can mitigate losses or remain profitable by implementing recovery policies—such as vehicle repossession, wage garnishment, and legal enforcement—and by charging relatively high interest rates that generate substantial cash flows early in the loan term. As a result, default-based measures alone may not fully capture the economic consequences of credit decisions.

We construct two measures of post-approval performance. The first measure, *Default*, captures repayment performance after loan approval. It is a binary indicator equal to one if the borrower becomes delinquent at any point during the loan term. Delinquency is defined as a failure to make a scheduled monthly payment for more than 60 days past the due date. The 60-day threshold is chosen to exclude short payment delays that are often due to temporary oversight and are typically resolved promptly following phone calls or SMS reminders.

The second measure, the *Loan Profit Ratio*, follows the approach in Lu et al. (2023) and Jansen et al. (2025). We first calculate the profit associated with each transaction as the present value of all cash flows generated over the life of the loan net of the initial resources invested by the firm. Cash inflows include monthly loan payments, delinquency-related fees, and, in the event of default, net recovery proceeds from repossession, auction, and post-default enforcement

actions, after deducting associated costs. The initial investment reflects the firm’s economic exposure at origination and is measured as the sum of the vehicle cost and the net lending outlay, defined as the approved loan amount net of the borrower’s down payment and any upfront fees.

For loans that are fully repaid, profit equals the present value of all payments plus the vehicle-side margin minus the initial investment. For loans that default, profit is calculated as the present value of payments received up to default plus the discounted value of post-default recoveries, again net of the initial investment. The loan profit ratio is defined as total profit scaled by the initial investment.

For loans that remain outstanding at the end of the sample period, we compute expected profitability by combining realized cash flows with the expected present value of future payments and recoveries, based on historical default and recovery patterns by loan age. We report results using realized profitability for completed loans and expected profitability for ongoing loans.

Coarse-group information reliance index. To evaluate fairness in credit approval, we adopt two widely used criteria—*equalized opportunity* and *demographic parity* (Teodorescu et al., 2021; Chohlas-Wood et al., 2026). Both criteria share a common implication: conditional on the same underlying credit quality (or comparable risk), borrowers from different groups should face the same probability of receiving the positive outcome (loan approval), thereby avoiding systematic disadvantages imposed on groups that should otherwise have access to financial services (Feldman et al., 2015). In this framework, the key deviation of interest is not merely individual-level heterogeneity, but the systematic reliance of the approval rule on *coarse-group information*. When decision-makers (human or algorithmic) evaluate individual credit risk by placing excessive weight on information that is highly correlated with group identity or group-level averages—and substitute such information for an applicant’s true credit potential—the resulting pattern corresponds to the classic notion of statistical discrimination (Phelps, 1972; Schwab, 1986).

In credit approval settings with intelligent machine adoption, statistical discrimination often manifests through the specific mechanism of *coarse-group information reliance*. We highlight three main channels. First, intelligent machines are typically trained on historical approval and default data; if past decisions reflect systematic group-level disparities, algorithms that minimize overall prediction error may inherit and amplify reliance on coarse-group information (Fuster et al., 2022; Dobbie et al., 2021). Second, when optimizing average predictive accuracy, machine-learning models tend to exploit statistical features that are strongly correlated with group labels (e.g., group-average default rates or their proxies), thereby substituting coarse-group averages for individual credit potential and reinforcing group-based judgment (Lin and Viswanathan, 2016; Zhao and Wry, 2016). Third, relative to traditional face-to-face underwriting,

intelligent machines shift the decision process toward structured hard information, weakening soft-information and contextual correction channels and further increasing the weight placed on coarse-group information in credit assessment (Hu et al., 2025).

To systematically measure the intensity of approval rules' reliance on coarse-group information in our sample, we construct a *coarse-group information reliance index* (CGIR). The index is designed not to directly capture outcome disparities, but to quantify the extent to which credit approval relies on group-label-type information. When the approval rule more frequently and more strongly depends on coarse-group characteristics related to demographic and socioeconomic stratification, it is more likely to substitute group-average risk for individual credit potential, thereby reflecting a higher degree of statistical discrimination.

Specifically, in the first step, drawing on the literature on discrimination, bias, and statistical substitution in credit approval (Lin and Viswanathan, 2016; Zhao and Wry, 2016; Fuster et al., 2022; Lu et al., 2023; Dobbie et al., 2021; Hu et al., 2025), we select a set of variables that are likely to carry coarse-group information in practice and thus serve as potential channels of statistical discrimination. Each variable corresponds to a demographic or socioeconomic group label, including: *Client Workplace* (whether the applicant holds formal employment), *Client House Ownership* (whether the applicant owns a house), *Married Client* (equal to 1 for married applicants and 0 for unmarried, divorced, or widowed applicants), *Has Children* (whether the applicant has children), *Local Client* (equal to 1 if the applicant's registered-place prefecture-level city matches the prefecture-level city of the loan contract, and 0 otherwise), *Male Client* (equal to 1 for male applicants and 0 for female applicants), *PrimeAge Client* (equal to 1 for applicants aged 25–54 and 0 otherwise), *CollegePlus Client* (equal to 1 for applicants with a bachelor's degree or above), *Repeated Client* (whether the applicant purchased a car from the firm before the start of our sample), and *Log(Client Wage)* (the logarithm of monthly income). The classification of *PrimeAge Client* follows the OECD definition of prime working age. Together, these variables capture applicants' socioeconomic status, family structure, local attachment, and prior relationship with the firm; in the absence of individualized soft information, they can be readily used as proxies for coarse-group-level risk.

In the second step, to verify that these coarse-group variables indeed embed empirically relevant group-level differences in the firm's historical practice, we conduct a validation exercise using the firm's historical data from the ten years prior to the start of our sample. In unreported analyses, we examine whether loans associated with different coarse-group labels exhibit systematically higher approval probabilities or different loan performance, conditional on observable hard information. This step does not equate historical differences with discrimination; rather, it establishes that these variables plausibly generate stable group-average signals and thus can serve as inputs for statistical substitution (Fuster et al., 2022; Dobbie et al., 2021).

In the third step, we compare the reliance on coarse-group information before and after FinTech adoption. Specifically, we train the same prediction models separately in the pre-adoption and post-adoption samples and rank variables by model-based feature importance to identify the features that are most influential for approval and risk prediction in each period.⁵ Table A3 reports the top-ranked features in each period. We find that after FinTech adoption, variables highly related to coarse-group labels become salient and consistently appear among the most important predictors, in particular *Client House Ownership*, *Client Workplace*, *Married Client*, *Local Client*, *Log(Client Wage)*, and *Has Children*. Accordingly, we focus on these six variables to measure coarse-group information reliance in the approval process.

Finally, to aggregate multidimensional coarse-group information reliance into a single index suitable for regression analysis, we normalize the six variables (mapping continuous variables to the 0–1 range while keeping binary variables as 0/1) and apply principal component analysis (PCA). We construct CGIR using the first principal component. A higher CGIR indicates stronger reliance on statistical information tied to group labels and a greater tendency to substitute group-average risk for individual credit potential, reflecting a higher degree of statistical discrimination. In the empirical analysis, we use this index as the core measure of coarse-group information reliance and examine its dynamic changes around FinTech adoption and its implications for credit decisions and outcomes.

C. Summary Statistics

Table A4 presents summary statistics on credit approval outcomes, reporting the number of approved applications by FinTech adoption period and by vehicle segment. The sample includes a total of 161,818 approved applications, of which 87.45% are commercial vehicle loans and 12.55% are passenger vehicle loans. Application volume exhibits a general upward trend over time, peaking in 2020 and 2021. Approvals issued after FinTech adoption account for 79.38% of the commercial vehicle subsample and 73.74% of the passenger vehicle subsample.

Figure 3 plots the distribution of FinTech credit scores for passenger vehicle applicants, including both approved and rejected applications. The distribution is right-skewed, with a substantial mass of observations concentrated around the cutoff score of 475. Based on this threshold, 42.67% of applications fall below the cutoff and are rejected, while the remaining

⁵Specifically, we implement two model classes that align with the firm’s scoring practice: (i) logistic regression and (ii) XGBoost. In each period, we train the models using the same feature set and the same data preprocessing pipeline (including identical missing-value handling and encoding rules). For logistic regression, we estimate an L2-regularized model with the regularization strength selected via cross-validation. For XGBoost, we use the standard gradient-boosted tree classifier and tune key hyperparameters (e.g., maximum depth, learning rate, and number of trees) via cross-validation within each period. Feature importance is computed from the fitted models (e.g., gain-based importance for XGBoost), and the ranking is reported for each period.

57.33% are approved. Figure 4 further depicts loan approval rates across score bins. Consistent with the scoring design, approval rates increase monotonically with the FinTech score. In particular, for applicants with scores above 600, the approval rate reaches 100%, indicating full alignment between credit officers’ decisions and the FinTech system’s recommendations for high-scoring applicants.

V. Dynamics of the Human Role in Human–Machine Collaboration

This section examines the dynamic role of human decision-makers in human–machine collaboration. We begin by providing an overall assessment of the credit effects of FinTech score adoption, focusing on its implications for loan profitability and fairness. We then turn to a subsample of rescued loans—applications initially rejected by the machine but subsequently approved by credit officers—to study how human judgment intervenes in machine-based decision-making. Finally, we investigate the dynamic evolution of the human role in the rescue process, with particular attention to how the effectiveness and behavioral patterns of human intervention change over time following FinTech adoption.

A. Credit Effects of FinTech Score Adoption

To examine the impact of human–machine collaboration on credit outcomes following the introduction of FinTech tools, our empirical specification employs a staggered difference-in-differences design using both commercial and passenger vehicle samples. This approach is motivated by the fact that the FinTech scoring system was rolled out at different points in time across the two vehicle segments. Exploiting this staggered adoption, we estimate the following regression model:

$$Y_{k,t} = \beta_0 + \beta_1 \text{Treat}_{v,t} + \mathbf{X}_{k,t} + \phi_{p,t} + \delta_{j,t} + \mu_{a,t} + \varepsilon_{k,t}, \quad (1)$$

where k indexes the transaction, v indexes the vehicle type (commercial or passenger), p indexes the province in which the contract is signed, j indexes the credit officer assigned to transaction k , a indexes the dealership (agent), and t indexes the year. The outcome variable $Y_{k,t}$ includes the default indicator (*Default*), the loan profit ratio (*Loan Profit Ratio*), and the coarse-group information reliance index (*CGIR*). The treatment variable $\text{Treat}_{v,t}$ captures the exogenous introduction of the FinTech credit scoring system. For commercial vehicles, $\text{Treat}_{v,t}$ equals one for applications submitted on or after March 2020 and zero otherwise. For passenger vehicles, the indicator equals one for applications submitted from July 2020 onward and zero otherwise.

The control vector $\mathbf{X}_{k,t}$ includes contract-, vehicle-, and applicant-level characteristics.⁶ The specification includes province-by-year fixed effects ($\phi_{p,t}$), credit-officer-by-year fixed effects ($\delta_{j,t}$), and dealership-by-year fixed effects ($\mu_{a,t}$).⁷ Standard errors are two-way clustered at the credit-officer and year levels.

Table 1 presents the estimates of the impact of FinTech credit scoring adoption on loan performance outcomes. In Columns (1)–(3), the treatment coefficient on the default indicator is negative and highly significant, indicating that loans originated after FinTech adoption exhibit substantially lower default rates. The effect remains robust when restricting the sample to uncensored loans—defined as loans whose contractual maturity does not extend beyond the end of the sample period—as well as to completed loans that fully terminate within the sample window. Columns (4)–(6) examine loan profitability. The magnitude of the effect is similar for uncensored and completed loans, suggesting that the profitability gains are not driven by sample truncation or differential censoring. Taken together, these results indicate that FinTech credit scoring improves ex post loan performance by both reducing default risk and enhancing economic returns. Table A5 further examines the effect of FinTech adoption on early-stage delinquency using default indicators measured at different horizons. Across specifications, we find no statistically significant effects on short-term default outcomes (e.g., within the first 6 to 36 months after origination). This pattern suggests that the overall decline in default documented in Table 1 is not driven by changes in early repayment behavior, but instead reflects improvements in longer-term loan performance.

Table 2 examines the fairness implications of FinTech adoption. Column (1) reports the effect on the composite CGIR measure. The estimated treatment coefficient is positive and highly significant, indicating that the introduction of algorithmic credit scoring leads to a marked increase in reliance on coarse-group information in credit approval decisions. This finding suggests that, following FinTech adoption, approval rules place greater weight on group-level statistical signals rather than solely on fine-grained, individual-specific information. Columns (2) through (7) decompose this aggregate effect by examining the individual components of

⁶Contract-level controls comprise the logarithm of contract duration in months ($\log(\text{Contract Duration})$), the ratio of the down payment to the vehicle’s assessed value (*Down Payment Ratio*), the logarithm of the total financing amount ($\log(\text{Contract Financing})$), and the annualized contract interest rate (*Contract Interest Rate*). Vehicle-level controls include an indicator for whether the vehicle is new (*New Vehicle*), an indicator for business registration status (*Business Use*), and a categorical variable for vehicle brand type (*Vehicle Brand*), coded as 0 for domestic brands, 1 for Sino–foreign joint ventures, and 2 for foreign brands. Applicant-level controls include *Client House Ownership*, *Client Workplace*, *Married Client*, *Local Client*, $\log(\text{Client Wage})$, and *Has Children*. In regressions where the outcome variable is *CGIR*, applicant-level controls are excluded to avoid mechanical overlap with the dependent variable.

⁷We do not additionally include vehicle-type fixed effects because officers are segmented by vehicle line (i.e., commercial-vehicle and passenger-vehicle applications are handled by different sets of officers), so the officer-by-year fixed effects already absorb vehicle-type-specific heterogeneity, while dealerships (agents) can serve both vehicle segments and are therefore shared across the two samples.

CGIR. We find that FinTech adoption significantly increases reliance on $\log(\text{Client Wage})$ and *Client Workplace*, as well as on *Client House Ownership*, *Married Client*, and *Has Children*. In contrast, the coefficient on *Local Client* is negative and statistically significant, indicating a reduced role of local residency in approval decisions after FinTech adoption, consistent with (Chen et al., 2022).

We further test whether heightened reliance on coarse-group information is associated with changes in loan contract terms in Table A6. We find that after the introduction of FinTech, borrowers with stronger coarse-group information are associated with higher loan-to-value (LTV) ratios, implying more permissive financing terms. By contrast, we find no statistically significant effect on contract interest rates.

Taken together, these results highlight a clear trade-off between efficiency and fairness following FinTech adoption. Algorithmic credit scoring improves credit performance by reducing default risk and increasing loan profitability, while simultaneously increasing reliance on coarse-group information in approval decisions and contract terms, thereby shifting credit allocation toward group-based statistical signals rather than purely individual-specific assessments.

Our staggered DiD identification relies on the absence of differential pre-trends across vehicle segments around the (segment-specific) FinTech rollout. Figure A4 presents an event-study version. Across outcomes, the estimated coefficients in the pre-adoption periods are small and statistically indistinguishable from zero, providing no evidence of systematic pre-trends prior to adoption. In contrast, the post-adoption coefficients display clear and persistent shifts: default declines while coarse-group information reliance rises, consistent with the baseline DiD estimates. As an additional falsification exercise, Table A7 assigns fictitious adoption dates that do not correspond to any policy or technology change. The placebo treatment effects are close to zero and statistically insignificant for default, loan profitability, and CGIR, mitigating concerns that our findings are driven by spurious time patterns or mechanical re-timing of the treatment indicator.

Because of substantial structural and behavioral differences between commercial and passenger vehicle segments, pooling the two samples—even with officer fixed effects—may raise comparability concerns. Moreover, the passenger vehicle segment offers richer information and a larger sample. To address these issues, we complement the staggered DiD analysis with a single DiD design focusing exclusively on the passenger vehicle sample. As detailed in Appendix B, we partition passenger vehicle loans into treated and control groups based on the level of financial development of the contract province. The results, reported in Tables A8, A9, and A10, are qualitatively similar to the main findings and confirm the robustness of our conclusions.

B. *The Human Role in Human–Machine Collaboration*

In this section, we examine human–machine collaboration by focusing on how human decision-makers respond to machine-generated recommendations. Specifically, we compare the performance and potential biases of loan applications that are initially rejected by the machine but subsequently approved by human auditors (“rescued” cases) with those that are jointly approved by both the machine and the human.

To assess the impact of human intervention, we estimate the following regression on a comparable sample of approved loan applications:

$$Y_{k,t} = \gamma_0 + \gamma_1 \text{Rescued}_{k,t} + \mathbf{X}_{k,t} + \phi_{p,t} + \delta_{j,t} + \mu_{a,t} + \varepsilon_{k,t}, \quad (2)$$

where, $\text{Rescued}_{k,t}$ is an indicator equal to one if an application is rejected by the machine but approved by the human auditor, and zero if it is approved by both. The control variables and fixed effects mirror those in Equation (1), and standard errors are clustered in the same manner.

One concern is that rescued and non-rescued applications may differ systematically in underlying credit risk, as machine scores are generated from high-dimensional applicant characteristics. To mitigate this concern, we restrict attention to a subsample of observations with machine scores in a narrow neighborhood around the assignment cutoff of 475 points. This design follows the intuition of a regression discontinuity framework, which facilitates comparison among applicants with similar predicted risk.⁸

Table 3 reports the regression results. Panel A focuses on loan outcomes. Column (1) shows that the coefficient on rescued for the default indicator is statistically insignificant, indicating that rescued loans do not exhibit a higher default probability relative to non-rescued loans. Column (2) reports the effect on the loan profit ratio and likewise finds no statistically significant difference, suggesting that human intervention does not reduce loan profitability. In contrast, Column (3) examines the coarse-group information reliance index and finds a negative and marginally significant coefficient. This result indicates that rescued cases rely less on group-level characteristics in approval decisions, highlighting the important role of credit officers in reducing dependence on coarse statistical proxies through human overrides.

Panel B further decomposes the decline in CGIR to identify which underlying dimensions

⁸Specifically, we determine the optimal bandwidth using the mean squared error–optimal bandwidth (MBW) criterion and retain observations within the interval $[475 - \text{MBW}, 475 + \text{MBW}]$ as a balanced comparison sample. Within this local window, we estimate the following local linear specification:

$$Y_{k,t} = \omega_0 + \omega_1 D_{k,t} + f(z_{k,t} - Z, D_{k,t}) + \mathbf{X}_{k,t} + \phi_{p,t} + \delta_{j,t} + \mu_{a,t} + \varepsilon_{k,t},$$

where $z_{k,t}$ denotes the machine-generated credit score, $Z = 475$ is the cutoff, and $D_{k,t} = \mathbf{1}\{z_{k,t} > Z\}$ is an indicator for being above the threshold. The function $f(\cdot)$ flexibly captures the relationship between outcomes and the running variable on either side of the cutoff.

drive this effect. We find that rescued applications are associated with significantly lower observable income, and a lower likelihood of home ownership and being married. These patterns suggest that human auditors tend to override machine rejections for applicants with weaker observable socioeconomic characteristics. At the same time, rescued borrowers are significantly more likely to be local clients, indicating that human auditors may incorporate context-specific or relationship-based information that is not fully captured by the machine scoring model. There is no statistically significant difference in formal employment status or parental status.

Taken together, the results in Panels A and B suggest that human auditors play a complementary role in human–machine collaboration. By selectively rescuing applicants with weaker observable characteristics but comparable repayment performance, human intervention mitigates excessive reliance on group-level predictors without increasing default risk or reducing profitability. These findings highlight how human judgment can partially offset the statistical discrimination tendencies inherent in automated credit scoring systems. This evidence is also consistent with (Lu and Zhang, 2025), which emphasizes that deliberate human “rethinking” can improve efficiency and fairness simultaneously.

C. *Dynamics of The Human Role in Human–Machine Collaboration*

To further examine how the unique value of credit officers in rescuing machine-rejected applications evolves over time, we divide the post–FinTech period for the passenger vehicle sample into two equal-length phases: an early period (July 2020 to March 2022) and a late period (April 2022 to December 2023), each spanning 21 months. This time window is sufficiently long to allow credit officers to repeatedly interact with the algorithm and observe outcome feedback, facilitating learning, adaptation, and the gradual formation of stable decision heuristics and behavioral norms, as emphasized in the psychology and organizational learning literature (March, 1991; Argote, 2012).

Figure 5 provides an overall depiction of final approval rates across different score bins. The figure shows that in the later period, applications rejected by the machine exhibit a further decline in final approval rates, while applications approved by the machine display a marked increase in approval rates. This pattern indicates that credit officers’ decisions increasingly converge toward machine recommendations, reflecting a strengthening alignment between human judgment and algorithmic outputs. The emergence of such alignment may be driven by two distinct mechanisms. The first is a *trust learning channel*, whereby credit officers gradually update their beliefs about algorithmic competence through repeated exposure to predictive performance and ex post feedback, leading them to shift from initial skepticism toward calibrated reliance and trust in machine recommendations. The second is a *career-concern channel*, under which credit officers initially intervene more actively due to replacement risk and reputational uncertainty

following algorithm adoption, but subsequently reduce effort and rely more heavily on the algorithm as its performance stabilizes and organizational roles become more clearly defined. Distinguishing between these two mechanisms is essential for understanding the dynamic evolution of the human role in human–machine collaboration, and constitutes a central focus of the analysis that follows.

Figure 6 further illustrates the distribution of machine scores for rescued applications—those rejected by the algorithm but ultimately approved by credit officers—across the two post-adoption phases. The shape of the distribution changes markedly over time. In the early phase, rescued applications are widely dispersed across the score spectrum, with a concentration in the 350–400 range and an approximately normal shape. By contrast, in the later phase the distribution shifts sharply toward the algorithmic cutoff of 475, with most rescued applications clustered in the 425–475 range and a pronounced reduction in low-score rescues (i.e., those far from the cutoff). Consistent with this shift, the mean score of rescued applications is substantially higher in the later phase than in the early phase (414.11 versus 383.88), indicating that human intervention increasingly targets applications that are marginally rejected by the machine.

This distributional shift is corroborated by the regression evidence in Table A11. Specifically, the number of rescued applications in the later period rises significantly within increasingly narrow neighborhoods within narrow score windows below the cutoff, corresponding to distances of 10, 20, and 30 points from 475. Moreover, using a continuous measure, the distance between rescued scores and the cutoff of 475 decreases significantly over time. Together, these results indicate that human intervention becomes progressively concentrated near the algorithmic decision boundary.

The dynamic change in the score distribution of rescued applications can be rationalized within the two mechanisms discussed above; however, the observed patterns do not allow us to cleanly distinguish between them. On the one hand, under the trust learning channel, as credit officers gain experience with the algorithm and become more confident in its overall predictive accuracy, they may strategically focus their attention on applications near the cutoff, where algorithmic uncertainty is greatest and human judgment is potentially most valuable. On the other hand, under the career-concern channel, concentrating rescues near the cutoff may also reflect defensive behavior, as intervening on marginal cases is easier to justify *ex post* and entails lower reputational risk than overriding the algorithm for clearly low-scoring applicants. Because both mechanisms generate similar empirical implications for the distribution of rescued scores, the evidence here is consistent with either channel, and disentangling their relative importance remains an open question for further analysis.

The above results jointly show that, as human–machine collaboration evolves, human intervention changes not only in intensity but also in its locus, gradually converging toward the

algorithm's decision margin. This pattern provides key evidence on the dynamic evolution of the human role in human-machine collaboration.

Beyond behavioral focus, Table 4 further examines the dynamic evolution of outcomes for rescued loans. For default, the coefficient on the interaction term is negative and highly significant, indicating that rescued loans in the later period exhibit substantially lower default probabilities relative to rescued loans in the early period. This pattern suggests that, as experience with the algorithm accumulates, credit officers increasingly rescue applications with more favorable risk profiles. Consistent with this interpretation, the positive coefficient on Rescued implies that in the early period, rescued loans are associated with higher default risk, an effect that is significantly attenuated in the later stage. In contrast, for the loan profit ratio, neither the main effects nor the interaction term are statistically significant. This finding indicates that the dynamic adjustment in human intervention primarily operates along the risk dimension rather than through changes in economic returns.

As for the fairness dimension. In the early period, rescued cases rely less on coarse-group information in approval decisions. However, this reduction in coarse-group information reliance weakens over time. In the later period, rescued loans increasingly resemble non-rescued loans in terms of their dependence on group-level characteristics. Specifically, In the early period, rescued loans are associated with significantly lower observable income, lower rates of home ownership, and a lower likelihood of being married. These patterns suggest that credit officers initially tend to override machine rejections for applicants with weaker observable socioeconomic characteristics. While In the later period, rescued loans are associated with significantly higher income levels and higher home ownership rates, indicating that human intervention increasingly concentrates on applicants with stronger observable characteristics who are nevertheless close to the machine decision boundary.

Overall, Table 4 shows that human intervention primarily operates through reducing statistical discrimination in the early period, while increasingly working through the screening of default risk in the later period. This dynamic pattern can also be explained by these two distinct mechanisms.

Under the *trust learning channel*, in the early stage following system adoption, loan officers have not yet fully developed trust in the algorithm's risk assessment capability. As a result, human intervention is more likely to systematically deviate from machine decisions, particularly by correcting potential group-level biases embedded in algorithmic outputs, thereby primarily improving fairness. As loan officers accumulate experience and learn about the predictive accuracy of the algorithm over time, their judgments gradually converge toward the machine-defined risk boundary. Consequently, human intervention shifts from broad corrective actions toward marginal adjustments, and in the later period operates mainly through lowering default

risk rather than correcting group disparities.

Under the *career-concern channel*, the dynamic pattern arises from changes in incentives and responsibility allocation. In the early stage of system introduction, the algorithm has not yet been fully embedded into the organization’s performance evaluation and accountability framework. Loan officers therefore face stronger career concerns and greater uncertainty, which incentivize them to engage in observable and differentiated interventions to demonstrate their judgment value. Such interventions are more likely to take the form of correcting group-level disparities. As the algorithm gradually becomes the standardized decision benchmark within the organization, individual officers’ responsibility boundaries and incentive structures stabilize. Human intervention correspondingly contracts and becomes concentrated on cases close to the machine’s decision margin, leading in the later period to improvements in loan outcomes primarily through risk control.

VI. Channel Analysis: Trust Learning versus Career Concerns

In this section, we distinguish between the trust learning channel and the career-concern channel along four complementary dimensions, aiming to assess which mechanism primarily drives the observed dynamic patterns. Specifically, we examine (i) textual characteristics of approval opinions, (ii) loan officers’ allocation of approval time as a proxy for effort and attention allocation, (iii) loan officer behavior during promotion periods, when performance evaluation pressure and career concerns are particularly salient, and (iv) heterogeneity across officers and loans.

A. Textual Analysis of Approval Opinions

Approval opinion texts associated with each loan constitute an important data advantage of this study. Compared with analyses that rely solely on approval outcomes or numerical variables, approval opinions provide a window—albeit an imperfect one—into the reasoning frameworks, informational emphasis, and expressive choices employed by loan officers in the decision-making process, thereby offering finer-grained evidence on human behavior under human–AI collaboration. Importantly, approval opinions should not be interpreted as a direct reflection of loan officers’ true psychological beliefs. Rather, they capture the judgments that officers choose to articulate under organizational constraints and institutional settings. However, precisely because the data provider’s internal guidelines do not impose strict requirements on the content, length, or format of approval opinions, these texts plausibly reflect meaningful variation

in written decision-making behavior across different stages and decision contexts.

More importantly, approval opinion texts provide a novel dimension for distinguishing between the two mechanisms proposed above. From a *trust learning* perspective, as loan officers gradually develop confidence in the algorithm's risk assessment capabilities over time, their approval opinions are expected to place greater emphasis on concrete facts, verifiable information, and explicit risk features, resulting in more focused and technical language. This tendency should be especially pronounced for marginal cases with credit scores close to the algorithmic cutoff, where algorithmic predictions are inherently more uncertain and loan officers are more likely to supplement machine output with substantive human judgment. In such cases, approval opinions are therefore expected to be more detailed, fact-based, and specific.

In contrast, under the *career concern* mechanism, as the algorithm increasingly becomes the standardized decision benchmark within the organization and serves as the primary anchor for accountability, loan officers may have stronger incentives to rely on reusable, highly generalized, and ex post defensible language in order to reduce personal responsibility exposure. In this setting, the linguistic style of approval opinions is more likely to exhibit increased vagueness and greater degrees of standardization. Notably, for marginal cases near the cutoff—where decisions can be more readily justified by deferring to the algorithm—loan officers face weaker incentives to provide extensive justification, leading approval opinions to become comparatively shorter and coarser in their articulation.

To empirically distinguish between the two mechanisms, we begin by examining approval opinion texts and characterizing loan officers' written decision behavior. We combine large language model-based text analysis with rule-based methods to construct four linguistic features that capture loan officers' behavior as reflected in approval opinions.

(i) Log(Text Length). Text length is defined as the total number of Chinese characters contained in an approval opinion. We use the logarithm of text length to reduce skewness in the distribution. This measure captures the baseline level of documentation effort at the textual level, reflecting the overall extent of written input provided by the loan officer, while remaining agnostic about the informational content or specificity of the text.

(ii) Information Density. Information Density measures the extent to which approval opinions contain verifiable, concrete information as opposed to generic or template-based language. A higher level of information density indicates that approval opinions include more specific facts, clearly defined attributes, and verifiable justifications, thereby reflecting a greater degree of substantive judgment and information integration by loan officers during the approval process.

Guided by credit approval practices and the content of approval texts, we construct three complementary subcomponents of information density. Each subcomponent is normalized

by the total number of tokens in the text to account for differences in text length. We then standardize and aggregate these components to obtain a composite Information Density index. Specifically, the three subcomponents are defined as follows. *Numeric Density* is defined as the frequency of numeric expressions in approval opinions, including both Arabic numerals and Chinese numerals, normalized by the total token count. Such expressions typically correspond to specific amounts, installment numbers, ratios, or quantities, are highly verifiable, and are unlikely to appear in generic or boilerplate language. *Factual Entity Density* measures the share of text that refers to concrete, case-specific facts, such as time, location, identity characteristics, verification processes, and verification outcomes. This component captures the extent to which loan officers explicitly document factual background information and objective verification results related to the applicant. *Feature-Referencing Density* captures explicit references to borrowers' economic characteristics and credit-related attributes, including income, liabilities, credit history, loan status, and risk identification outcomes. This component reflects the degree to which approval opinions directly engage with quantifiable risk-relevant features. Table A12 reports the complete keyword dictionaries used in constructing the Factual Entity Density and Feature-Referencing Density measures.

(iii) Fuzzy Dummy. The Fuzziness Dummy captures whether an approval opinion contains vague, highly generalized, and easily defensible language that lacks explicit factual justification. Such expressions typically do not refer to specific numerical values, verifiable facts, or well-defined borrower characteristics. Instead, they rely on summary or judgment-based phrasing to complete the approval rationale, thereby reducing individual decision-makers' accountability exposure.

To construct this measure, following the approach in (Baker et al., 2024), we employ a Sentence-BERT-based large language model and adopt a three-step procedure that combines dictionary screening, human annotation, and semantic model identification. First, based on credit approval practices and the content of approval texts, we manually compile a dictionary of frequently used fuzzy expressions (e.g., “basically compliant,” “overall acceptable,” “risk controllable”). The complete list of keywords is reported in Table A13. Using this dictionary, we identify sentences in approval opinions that contain at least one fuzzy expression and treat them as a set of candidate fuzzy sentences. Second, we randomly sample sentences from this candidate set for manual annotation. Each sentence is labeled as either semantically fuzzy/defensive or non-fuzzy based on whether it lacks verifiable justification, fails to reference concrete facts or features, or substitutes specific reasoning with generalized evaluative statements. To ensure annotation quality, we employ independent double coding and compute inter-annotator agreement. Discrepancies are resolved through discussion to obtain a final label. Third, we use the manually labeled sentences to fine-tune a Chinese Sentence-BERT model in a supervised

setting, allowing the model to learn shared semantic patterns associated with fuzzy expressions. The trained model is then applied to the full corpus of approval opinions at the sentence level. An approval opinion is assigned a value of one for the Fuzziness Dummy if it contains at least one sentence classified by the model as semantically fuzzy; otherwise, the variable is set to zero.

(iv) Textual Similarity. Textual similarity measures the extent to which approval opinions become standardized and repetitive within the same decision environment, capturing semantic similarity across different approval justifications. Unlike approaches based on word frequency or character-level overlap, we rely on text embeddings generated by a pre-trained large language model (LLM) to represent the semantic content of approval opinions, which allows us to capture higher-level semantic similarity. Compared with TF-IDF-based similarity measures, embedding-based similarity does not depend on the overlap of specific words or phrases, but instead leverages the compressed semantic representations learned by the pre-trained model. As a result, it is better suited to assessing whether approval opinions rely on substantively similar reasoning rather than merely sharing surface-level textual patterns.

Specifically, we treat each approval opinion as an independent document and use a pre-trained LLM to generate a corresponding semantic embedding vector. Within each decision environment—defined by loan officer \times year \times credit score bin—we compute the cosine similarity between all pairs of embedding vectors and take the average as the group-level textual similarity measure. This measure is then assigned to each loan application within the group. A higher value of textual similarity indicates that approval opinions within the same decision environment are semantically more homogeneous, reflecting more template-based and less individualized reasoning. Conversely, lower similarity suggests greater variation in expressed reasoning and a higher degree of case-specific judgment.

Table 5 characterizes the overall changes in approval opinion text following the introduction of the FinTech credit scoring system. Panel A shows that the adoption of FinTech itself leads to two salient effects. First, the overall length of approval opinions becomes significantly shorter, indicating that algorithmic support reduces loan officers’ baseline documentation effort. Second, the use of fuzzy expressions increases markedly, suggesting that approval opinions increasingly rely on generalized and reusable language.

Panel B further reveals that these changes are amplified over time. Compared with the early post-adoption period, approval opinions in the later period exhibit further reductions in text length and information density, accompanied by higher usage of fuzzy expressions and greater textual similarity. This evolution in textual behavior indicates that, as the algorithm gradually becomes the dominant decision benchmark within the organization, loan officers do not reinforce their independent judgment by incorporating more factual or feature-based explanations. Instead, they systematically shift toward more concise, vague, and standardized language. When decision

responsibility can be partially delegated to the algorithm, such reusable and ex post defensible expressions help reduce individual accountability exposure. Taken together, these patterns point to career concerns—rather than trust learning—as the primary mechanism underlying the observed textual changes in the later stage of human–AI collaboration.

Table 6 further examines the dynamic patterns of approval opinion text for marginal cases around the algorithmic cutoff. Panel A shows that, in the later period, applications with credit scores close to the cutoff are associated with shorter approval opinions, lower information density, and higher usage of fuzzy expressions. Panel B extends the discrete cutoff analysis by measuring proximity to the cutoff using the continuous score distance, yielding consistent results. Table A14 provides additional robustness checks by adopting alternative definitions of the cutoff neighborhood, all of which lead to similar findings.

The observed contraction and increased fuzziness of approval opinions for cutoff-adjacent cases in the later period further support the career concern mechanism. Under a trust learning framework, cases near the cutoff—where algorithmic predictions are most uncertain—should be precisely those in which human judgment plays a greater complementary role, leading to more detailed and fact-based textual explanations. Instead, we observe the opposite pattern. As the algorithm becomes a reliable decision benchmark and a central anchor for accountability, loan officers reduce textual effort even in the most contentious marginal cases, relying more heavily on reusable and easily defensible language to mitigate personal responsibility. This evidence further reinforces the dominant role of career concerns in shaping loan officers’ behavior in the later stage of human–AI collaboration.

B. Approval Time Allocation and Effort Provision

This dimension can also provide a sharp distinction between the trust learning channel and the career-concern channel. If human intervention primarily operates through *trust learning*, then as loan officers become more familiar with and confident in the algorithm’s predictive ability, their limited attention and effort should be increasingly reallocated toward applications for which the machine’s judgment is most uncertain—namely, those with credit scores close to the algorithmic cutoff. Under this mechanism, one would therefore expect greater time investment in cutoff-adjacent cases in the later period. In contrast, if human intervention is mainly driven by *career concerns*, then as the system matures and algorithmic decisions become the organization’s standard benchmark, loan officers may rely more heavily on machine outputs to mitigate accountability risk. In this case, officers would devote less time to marginal applications that are rejected by the algorithm but potentially overturned by humans, as such decisions are easier to justify ex post by referencing algorithmic recommendations.

Our data include transaction-level audit time for all passenger vehicle loan applications,

which allows us to directly examine how loan officers allocate time across different stages of system adoption. Specifically, we measure the average time spent per application in minutes and construct the log-transformed variable $\text{Log}(\text{Audit Time})$. Since junior screening officers are primarily responsible for verifying the authenticity of submitted materials, while discretionary judgment is concentrated in the final approval stage, our analysis in this subsection focuses on senior review officers.

Figure 7 presents the overall evolution of audit time following the adoption of the FinTech credit scoring system and reveals four key patterns. First, average audit time per application declines substantially from 68.29 minutes before adoption to 50.16 minutes after adoption, consistent with prior evidence that algorithmic assistance improves screening efficiency (Fuster et al., 2019). Second, in the post-adoption period, applications rejected by the algorithm (low-score cases) receive significantly more audit time than those approved by the algorithm (high-score cases), and the time spent on low-score applications even exceeds the pre-adoption average, indicating continued human scrutiny in machine-rejected cases. Third, audit time in the early post-adoption period is substantially higher than in the late period (approximately 67.27 versus 31.16 minutes) and remains close to the pre-adoption level, suggesting a nontrivial adjustment phase in early human-machine collaboration, during which efficiency gains do not materialize immediately. Fourth, in the late post-adoption period, audit time declines markedly for both high- and low-score applications, reflecting a further contraction and stabilization of the approval process. Taken together, these patterns indicate substantial time reallocation by loan officers as the system evolves.

To further distinguish between the two mechanisms, Figure 8 plots audit time across credit score bins separately for the early and late post-adoption periods. Relative to the early period, loan officers in the late period spend significantly less time on applications near the algorithmic cutoff, with the contrast being most pronounced precisely around the cutoff. This pattern runs counter to the prediction of the trust learning channel, which would imply increased attention to marginal cases as experience accumulates. Instead, it is more consistent with the career-concern channel: in the later stage of system operation, loan officers appear to rely more heavily on algorithmic recommendations to limit accountability exposure, reducing additional scrutiny of marginal cases that are rejected by the machine but potentially approved by humans. As a result, human intervention in the late period takes on a more defensive and procedural form rather than involving intensive case-by-case evaluation.

Table 7 provides regression-based evidence that corroborates the graphical findings. In particular, Panel B shows that, in the late period, applications close to the algorithmic cutoff do not receive more audit time and instead experience a statistically significant contraction in time allocation. These results offer further support for the career-concern channel in explaining the

evolution of human intervention in the later stage of human–machine collaboration.

C. Loan Officer Behavior under Promotion Pressure

Evidence from approval time allocation and approval-text characteristics in the previous sections indicates that behavioral changes in the later stage of human–machine collaboration are primarily driven by career concerns rather than trust learning. To further substantiate this mechanism, this section examines loan officer behavior in an environment where career incentives are exogenously intensified, namely during promotion periods.

Promotion periods refer to specific time windows following the launch of certain loan products, typically lasting between two and six months. Relative to regular periods, promotion periods are characterized by stronger performance pressure, more intensive internal rankings, and closer peer comparison. In this environment, individual loan officers’ performance becomes more observable, more frequently evaluated, and more tightly linked to career advancement, thereby substantially amplifying career concerns. Under the career-concern mechanism, heightened performance pressure strengthens loan officers’ incentives to actively demonstrate their judgment ability. In such settings, officers are more likely to deviate from algorithmic recommendations and intervene in or adjust machine decisions in order to signal their value added and perceived irreplaceability.

Figure 9 compares the credit score distributions of rescued applications during promotion and non-promotion periods. Subfigure (a) shows that, relative to non-promotion periods, rescued loans during promotion periods exhibit a pronounced leftward shift in the credit score distribution, with a significantly lower mean score. This pattern indicates that loan officers are more willing to extend their discretion beyond algorithmic boundaries during promotion periods, actively intervening in applications with weaker credit quality. In a high-stakes performance environment, officers have stronger incentives to engage in observable, discretionary actions that demonstrate their contribution, even at the cost of higher risk exposure.

Subfigure (b) further shows that this leftward shift becomes more pronounced in the late post-adoption period. Combined with earlier findings that, in regular periods, human intervention in the later stage tends to converge toward the algorithmic cutoff, promotion periods introduce a clear countervailing force. Rather than concentrating closer to the cutoff, rescued applications during promotion periods expand more aggressively into lower-score regions where algorithmic rejection is stronger. This contrast highlights how promotion-induced incentives alter behavior in the late stage of human–machine collaboration: when the algorithm has become a stable decision benchmark, loan officers typically adopt a defensive, easily justifiable strategy by intervening only at the margin; however, during promotion periods, additional performance incentives induce officers to deviate from this “safe strategy” and engage in more salient discretionary interventions

to enhance the visibility of their contribution. As a result, the usual convergence toward the cutoff observed in the later stage is partially offset or even reversed. These distributional patterns are systematically confirmed by the regression results reported in Table 8.

Table 9 further examines the performance consequences of rescued loans during promotion periods. Overall, the results indicate that promotion-period interventions are more explicitly performance-oriented. On the one hand, rescued loans during promotion periods are associated with lower default risk and, in some specifications, higher profitability, suggesting that loan officers actively deploy discretion to expand lending volume and improve short-term performance under strong incentive pressure. On the other hand, this performance-driven intervention comes with a fairness cost: discrimination measures increase for promotion-period rescued loans, indicating a temporary exacerbation of inequality. From a dynamic perspective, the increase in discrimination attenuates in the later post-adoption period. Consistent with the earlier distributional evidence, this pattern reflects changes in the composition of rescued loans over time. In the later stage, promotion-period intervention becomes increasingly concentrated among lower-score applicants, which weakens the role of observable group-level differences and leads discrimination measures to decline. Taken together, these results show that promotion periods induce loan officers to intensify discretionary intervention to improve default and profitability outcomes, but at the cost of temporary and structural fairness trade-offs. This evidence aligns closely with the career-concern mechanism identified earlier and reinforces the central role of performance pressure in shaping human behavior under human-machine collaboration.

D. Heterogeneity Analysis

This subsection provides additional mechanism evidence by examining whether the late-stage contraction of human effort and the increased standardization of textual justifications are systematically stronger in environments where career concerns and accountability pressures should matter most. The key prediction under the career-concern channel is that, as the algorithm becomes an organizational benchmark, officers reduce discretionary effort particularly for marginal, contestable cases (i.e., those near the cutoff), and this tendency is amplified when (i) reputational stakes and replacement risk are higher, (ii) prior reputation is weaker, (iii) ability is more uncertain, or (iv) external uncertainty elevates the demand for defensible decision-making. In contrast, a pure trust-learning view would more naturally predict increased attention and richer justification for cutoff-adjacent cases in the later period.

We focus on rescued passenger-vehicle applications and examine five process-level outcomes that capture officers' effort provision and articulation: $\text{Log}(\text{Audit Time})$, $\text{Log}(\text{Text Length})$, Information Density, Fuzzy Dummy, and Textual Similarity. We then study whether the late-stage near-cutoff behavioral shift varies systematically across officer- and environment-level

characteristics that, under the career-concern framework, should affect accountability pressure and the incentives to produce defensible decisions. Specifically, we consider six heterogeneity measures. *Tenure* is the number of months an officer has been in the senior-approval position at the time of FinTech adoption, capturing where the officer is on the career ladder when the algorithm arrives. *Past Profit Ratio* and *Past CGIR* summarize the officer’s pre-FinTech track record (average loan profit ratio and average CGIR among pre-FinTech approvals), proxying for baseline reputation and decision style. *Vol(Past Profit Ratio)* and *Vol(Past CGIR)* are within-officer standard deviations of monthly pre-FinTech profit ratio and CGIR, capturing uncertainty about the officer’s ability or consistency. Finally, *COVID Exposure* is the city-year share of days classified as medium/high risk, capturing external uncertainty and operational stress that may increase the demand for ex post defensibility.

Table 10 reports the results. Across panels, a clear pattern emerges: the late-stage cutoff region is precisely where officers’ behavior becomes more “defensive”—less time-intensive, less substantively articulated, and more template-like—and this tendency is stronger in settings where career concerns and accountability pressures are expected to be more salient.

Panel A shows that the late-stage near-cutoff contraction is more pronounced along the tenure dimension. Officers with higher tenure exhibit a sharper reduction in audit time and written elaboration, with approval texts becoming shorter and less information-dense. At the same time, their justifications display greater use of vague phrasing and higher semantic similarity, consistent with a shift toward reusable and easily defensible language rather than case-specific reasoning.

Panel B indicates that the same late-stage near-cutoff pattern also varies with officers’ historical performance. Conditional on fixed effects and observables, officers’ pre-FinTech profitability track record predicts how strongly their late-stage near-cutoff behavior tilts toward reduced effort and more standardized articulation. Panel C provides complementary evidence from the fairness dimension: officers with a historically higher reliance on coarse-group information exhibit a stronger late-stage near-cutoff move toward shorter, less information-rich explanations and more standardized language. Taken together, Panels B and C suggest that the late-stage “algorithm-as-benchmark” environment disproportionately compresses documentation effort and encourages standardized justifications among officers whose historical decision patterns align more with coarse, defensible screening heuristics.

Panels D and E speak directly to the role of ability uncertainty. Panel D shows that when officers’ pre-FinTech profitability performance is more volatile, the late-stage near-cutoff environment is associated with a larger decline in time investment and substantive writing, accompanied by more fuzziness and greater text standardization. This is consistent with the notion that greater uncertainty about performance-relevant ability amplifies incentives to limit exposure

and rely on defensible, low-variance documentation strategies. By contrast, Panel E shows that volatility in pre-FinTech CGIR is comparatively less predictive—most notably, it does not translate into a statistically discernible change in Information Density. This divergence is informative: what appears most consequential for late-stage marginal-case behavior is uncertainty in performance-relevant outcomes (profitability), rather than dispersion in group-information reliance per se.

Finally, Panel F uses *COVID Exposure* as an external uncertainty shifter. Under heightened external uncertainty, approval texts near the cutoff in the late period become more defensive—more likely to contain vague language and more semantically standardized—even when changes in time spent and text length are less pronounced. This pattern aligns with accountability considerations: when the environment becomes more uncertain, officers’ written articulation shifts toward language that is easier to justify *ex post*, particularly for marginal cases where the algorithm provides a salient reference point.

Overall, the heterogeneity results reinforce the career-concern mechanism. The late-stage convergence toward algorithmic benchmarking is not uniform; it is systematically stronger for officers and environments where reputational stakes, performance uncertainty, and external stress plausibly raise the value of defensibility. Importantly, the direction of these patterns—reduced scrutiny and more standardized justification precisely around the cutoff in the later period—runs counter to a trust-learning account, which would predict greater attention and richer case-specific articulation where algorithmic uncertainty is highest.

VII. Further Analysis: Dynamics of Upward and Downward Human Overrides

In the preceding analysis, we focus primarily on applications that ultimately receive credit, namely Group A (machine approve + human approve) and Group B (machine reject + human approve). By comparing Group A with Group B, we document the effects of corrective human intervention in human–machine collaboration, where loan officers override algorithmic rejections and approve a subset of applications.

An important advantage of our data is that we also observe the full set of applications that ultimately fail to pass credit screening, including Group C (machine approve + human reject) and Group D (machine reject + human reject). This complete partition of applications allows us to further examine the role of human intervention from two additional and complementary perspectives. First, among all applications initially rejected by the machine, loan officers do not uniformly overturn algorithmic decisions but selectively implement upward overrides that result in final approval. By comparing Group B with Group D, we can evaluate the marginal impact of

upward overrides within the pool of low-score applications rejected by the algorithm. Second, among applications initially approved by the machine, loan officers retain veto power and implement downward overrides for a subset of cases, leading to final rejection. By comparing Group A with Group C, we further analyze the role of human discretion in risk control and cautious screening beyond algorithmic approval.

Before turning to the two sets of comparisons described above, Figure 10 provides a descriptive overview of the distribution of the four groups in the full sample and across different periods. Subfigure (a) shows that Group A constitutes the largest share of applications, indicating that machine and human decisions are aligned in the majority of cases. Group B also represents a sizable fraction of the sample, suggesting that rescued decisions are not a marginal phenomenon but rather an important component of human–machine collaboration. In contrast, Group C is relatively small in size; combined with Figure 4, these cases are primarily concentrated in the 475–600 score range. Group D still accounts for a non-trivial share, indicating that while human vetoes of machine decisions are relatively cautious, they are not absent.

Subfigure (b) further splits the sample into the early and late post-adoption periods and reveals informative changes in relative proportions. Among machine-approved applications (Group A + Group C), the share of downward overrides (Group C) declines markedly in the late period: in the early period, Group C accounts for approximately 8.7% of machine-approved applications ($432/(4535+432)$), whereas in the late period this share falls to about 6.5% ($331/(4770+331)$). Similarly, among machine-rejected applications (Group B + Group D), the share of upward overrides (Group B) also contracts over time. In the early period, Group B constitutes roughly 61.5% of machine-rejected applications ($2703/(2703+1696)$), while in the late period this proportion decreases to approximately 55.8% ($1861/(1861+1474)$).

Taken together, both downward and upward overrides become systematically less frequent in the late post-adoption period. This pattern is highly consistent with the career-concern mechanism driven by increased defensibility: as the algorithm becomes a widely accepted and institutionally legitimate decision benchmark, relying on machine outcomes itself becomes a safe and easily defensible strategy. Consequently, loan officers are less inclined to deviate from algorithmic decisions in either direction, leading to a contraction of human intervention over time.

A. *Dynamics of Upward Human Overrides*

This subsection focuses on applications that are initially rejected by the machine. Within this subsample, we study cases in which loan officers subsequently approve the application despite the machine’s rejection, which we refer to as *upward human overrides*. These cases correspond to Group B, while applications that remain rejected by both the machine and the

loan officer (Group D) serve as the relevant comparison group. To evaluate the performance consequences of upward human overrides, we require counterfactual measures of default risk and loan profitability for applications that would not have been approved absent human intervention. Following Fuster et al. (2019), we construct these counterfactual outcomes by leveraging ex post repayment behavior and model-based predictions.⁹

Table A15 presents the average effects of upward overrides. The results show that applications subject to upward overrides subsequently exhibit better loan performance, with lower default rates and higher profitability, indicating that loan officers are able to effectively correct machine rejections. Importantly, this improvement does not come at the expense of fairness. The CGIR measure declines on average for upward overrides, and the borrower-characteristics results further show a reduced reliance on observable group attributes such as formal employment status (*client workplace*) and local residency (*local client*). Taken together, these patterns suggest that when re-evaluating machine-rejected applications, loan officers rely less on coarse group-based proxies and instead engage in more individualized assessments.

Table 11 further examines how these effects evolve over time. In the late post-adoption period, upward overrides generate stronger improvements in loan outcomes, primarily by more effectively identifying applicants with lower ex post default risk. At the same time, CGIR increases in the later period, indicating a partial re-emergence of group-based dependence. This combination of stronger efficiency gains but higher group reliance is consistent with earlier evidence that human intervention increasingly concentrates near the algorithmic cutoff over time. As upward overrides become more focused on marginal cases, loan officers can more precisely select low-risk rejected applicants, thereby improving default performance, but may also revert to observable group attributes as defensible and expedient screening devices, leading to higher CGIR.

Overall, the two tables jointly show that upward human overrides in the machine-rejected region do not reflect a simple relaxation of standards. Instead, the decision logic evolves over time: in the early period, upward overrides emphasize correction and de-grouping, whereas

⁹Specifically, for approved loans, default outcomes and realized profits are directly observed. For applications that are rejected—either by the machine alone or jointly by the machine and the loan officer—we approximate their counterfactual repayment performance using predicted default probabilities and expected cash flows generated from a pre-trained risk model. The prediction model is estimated using an independent training sample in which repayment outcomes are fully observed, and it conditions on the same set of borrower, contract, and vehicle characteristics available at the time of decision. This approach allows us to place approved and rejected applications on a common outcome scale, thereby enabling a meaningful comparison of risk and profitability across override and non-override decisions. Loan profitability is computed as the expected present value of future cash flows net of the initial investment, incorporating predicted default probabilities and recovery rates in the event of default, consistent with the methodology in Fuster et al. (2019). By relying on counterfactual predictions rather than realized outcomes alone, this framework isolates the selection effects induced by human overrides and permits a clean assessment of whether upward overrides systematically improve or worsen loan performance relative to the machine’s baseline decision.

in the later period they increasingly reflect marginal risk screening accompanied by a partial return to group-based heuristics. This dynamic pattern aligns closely with the career concern and defensibility mechanism identified earlier. As the algorithm becomes a stable organizational benchmark, human intervention contracts and concentrates on more interpretable and defensible marginal cases, yielding more stable efficiency improvements while limiting further gains in fairness.

B. Dynamics of Downward Human Overrides

This subsection restricts attention to applications that are initially approved by the machine and examines loan officers' *downward human overrides* within the high-score approval pool. A downward override corresponds to Group C, with Group A serving as the comparison group. As in the analysis of upward overrides, the evaluation of default risk and loan profitability in this subsection is based on counterfactual outcome estimates constructed using ex ante observable information.

Table A16 reports the average effects of downward overrides. The results show that applications subject to downward override exhibit substantially worse counterfactual loan performance, with higher default risk and lower profitability. This pattern indicates that, even within the machine-approved region, loan officers are able to identify and screen out applications with unfavorable risk–return profiles. In this sense, downward overrides do not represent arbitrary interference with algorithmic decisions, but rather function as a form of risk backstop and quality control. Consistent with this interpretation, the borrower characteristics of downward-overridden applications are also systematically weaker. These borrowers tend to have lower wages and are more likely to be unmarried and without children, suggesting that loan officers continue to rely on observable individual characteristics to reassess risk even after a machine approval has been issued. Overall, the evidence in Table A16 implies that human vetoes within the machine-approved pool are, on average, effective in removing lower-quality approvals.

Table 12 further examines the dynamic evolution of these effects. In the late post-adoption period, applications subject to downward override are associated with even higher ex post default risk, indicating that human vetoes continue to play a role in risk containment. At the same time, however, these vetoed applications also display higher profitability, implying that late-stage downward overrides increasingly eliminate loans that would have generated relatively strong returns. In other words, the screening precision of downward overrides deteriorates over time: while default risk is reduced, this comes at the cost of forgoing profitable lending opportunities. Taken together, these findings suggest that downward human overrides become more conservative but also coarser as the system matures. Although loan officers remain effective at limiting downside risk, their veto decisions no longer dominate along both the risk and return

dimensions as they did earlier. This pattern is consistent with the broader evidence of behavioral convergence documented above. As the algorithm becomes an accepted and stable decision benchmark, human intervention increasingly reflects defensive risk control rather than finely calibrated risk–return optimization.

VIII. Conclusion

This paper studies how human behavior and collaborative value evolve over time in algorithm-assisted credit approval. Using transaction-level data from a nationwide in-house auto finance leasing firm in China, together with rich approval-opinion texts and transaction-level processing-time measures, we examine how loan officers interact with a machine-learning credit scoring system and how the efficiency–fairness profile of collaboration changes as the system matures.

We document three main findings. First, algorithmic scoring improves average performance, reducing default risk and improving economic outcomes, but it also increases reliance on coarse group-level proxies in approval decisions, indicating heightened statistical discrimination. Second, human discretion creates corrective value through rescued decisions: conditional on approval, rescued loans do not exhibit worse default or profitability outcomes than jointly approved loans, yet they display significantly lower coarse-group information reliance. Third, both the locus and the nature of human value shift over time. Rescued activity increasingly concentrates near the cutoff, and collaborative value reallocates from stronger fairness improvement early on toward stronger risk screening later, while profitability effects remain largely unchanged.

Mechanism evidence favors a career-concern interpretation over trust learning. Over time, approval-opinion texts become shorter, less information-dense, and more standardized, with greater reliance on vague and defensible language, especially near the cutoff in the later period. Audit time also contracts sharply over time, including for marginal cases, consistent with reduced effort provision as the algorithm becomes the organizational accountability benchmark. Promotion periods, which intensify evaluation pressure, revive more salient discretionary intervention and expand rescues into lower-score regions, further underscoring the role of incentives and principal–agent frictions.

Heterogeneity patterns reinforce this incentive-based interpretation. The late-stage, near-cutoff contraction in effort and articulation is most pronounced among officers whose career incentives and accountability pressures are plausibly stronger—those with shorter tenure, more volatile pre-FinTech performance, and operating under greater external uncertainty (e.g., higher COVID exposure)—consistent with a shift toward defensible, standardized behavior when reputational risk is salient.

We also leverage the full partition of applications into four machine–human decision groups.

Upward overrides within the machine-rejected pool improve selection and reduce coarse-group reliance on average, but over time they become more margin-focused and exhibit a partial return to group-based screening. Downward overrides within the machine-approved pool function as a risk backstop, yet their precision deteriorates as the system matures, increasingly sacrificing profitability to achieve conservative risk containment. Together, these results show that human discretion is not fixed: both its intensity and its effectiveness evolve as algorithmic decision support becomes institutionalized.

Our findings suggest concise implications for policy and organizational design. First, governance of hybrid decision systems should be dynamic: monitoring should track not only default outcomes but also profitability and systematic reliance on group proxies over time. Second, sustaining the value of human discretion requires incentive-compatible accountability and documentation regimes; otherwise, mature systems may induce defensive reliance on the algorithm and attenuated human effort. More broadly, the long-run benefits of human-machine collaboration depend as much on institutional design as on model accuracy.

REFERENCES

- Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo, 2022, Artificial intelligence and jobs: Evidence from online vacancies, *Journal of Labor Economics* 40, S293–S340.
- Adelino, Manuel, Kristopher Gerardi, and Barney Hartman-Glaser, 2019, Are lemons sold first? Dynamic signaling in the mortgage market, *Journal of Financial Economics* 132, 1–25.
- Agarwal, Sumit, and Itzhak Ben-David, 2018, Loan prospecting and the loss of soft information, *Journal of Financial Economics* 129, 608–628.
- Angwin, Julia, Jeff Larson, Surya Mattu, and Lauren Kirchner, 2022, Machine bias, in *Ethics of data and analytics*, 254–264 (Auerbach Publications).
- Argote, Linda, 2012, *Organizational learning: Creating, retaining and transferring knowledge* (Springer Science & Business Media).
- Autor, David H, and David Dorn, 2013, The growth of low-skill service jobs and the polarization of the US labor market, *American economic review* 103, 1553–1597.
- Bai, Bowen, Hengchen Dai, Dennis J Zhang, Lijian Jin, and Liang Liang, 2022, The impacts of algorithmic work assignment on fairness perceptions and productivity: Evidence from field experiments, *Manufacturing & Service Operations Management* 24, 3060–3078.
- Baker, Andrew C, David F Larcker, Charles G McClure, Durgesh Saraph, and Edward M Watts, 2024, Diversity washing, *Journal of Accounting Research* 62, 1661–1709.
- Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace, 2022, Consumer-lending discrimination in the FinTech era, *Journal of Financial Economics* 143, 30–56.
- Berg, Tobias, 2015, Playing the devil’s advocate: The causal effect of risk management on loan quality, *The Review of Financial Studies* 28, 3367–3406.
- Brennan, Michael J., Vojislav Maksimovic, and Josef Zechner, 1988, Vendor Financing, *The Journal of Finance* 43, 1127–1141.
- Brynjolfsson, Erik, and Tom Mitchell, 2017, What can machine learning do? Workforce implications, *Science* 358, 1530–1534.
- Buçinca, Zana, Maja Barbara Malaya, and Krzysztof Z Gajos, 2021, To trust or to think: cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making, *Proceedings of the ACM on Human-computer Interaction* 5, 1–21.
- Cao, Sean, Wei Jiang, Junbo Wang, and Baozhong Yang, 2024, From man vs. machine to man+ machine: The art and AI of stock analyses, *Journal of Financial Economics* 160, 103910.
- Castelo, Noah, Maarten W Bos, and Donald R Lehmann, 2019, Task-dependent algorithm aversion, *Journal of marketing research* 56, 809–825.
- Chalfin, Aaron, Oren Danieli, Andrew Hillis, Zubin Jelveh, Michael Luca, and Jens Ludwig, 2016, Productivity and Selection of Human Capital with Machine Learning, *American Economic Review: Papers & Proceedings* 106, 124–127.
- Chen, Yanting, Yingwei Dong, Jiayin Hu, and Yiping Huang, 2022, Does FinTech Reduce Human Biases? Evidence from Two Quasi-Experiments, *Evidence from Two Quasi-Experiments (December 9, 2022)* .
- Chod, Jiri, 2015, Trade Credit and Inventory Management, *Production and Operations Management* 24, 1560–1573.
- Chod, Jiri, Evgeny Lyandres, and Xuanming Yang, 2019, Trade Credit and Supplier Competition, *Journal of Financial Economics* 131, 484–505.
- Chohlas-Wood, Alex, Madison Coots, Henry Zhu, Emma Brunskill, and Sharad Goel, 2026, Learning to be fair: A consequentialist approach to equitable decision making, *Management Science* 72, 456–473.
- Choudhury, Prithwiraj, Evan Starr, and Rajshree Agarwal, 2020, Machine learning and human capital complementarities: Experimental evidence on bias mitigation, *Strategic Management Journal* 41, 1381–1411.
- Cowgill, Bo, 2018, Bias and productivity in humans and algorithms: Theory and evidence from resume screening, *Columbia Business School, Columbia University* 29.
- Cowgill, Bo, 2019, Bias and productivity in humans and machines, Technical report, Upjohn Institute Working Paper.
- Cowgill, Bo, and Catherine E Tucker, 2020, Algorithmic fairness and economics, *Columbia Business School Research Paper* .

- Daripa, Arup, and Jeffrey Nilsen, 2011, Ensuring Sales: A Theory of Inter-firm Credit, *American Economic Journal: Microeconomics* 3, 245–279.
- Davenport, Thomas, Abhijit Guha, Dhruv Grewal, and Timna Bressgott, 2020, How artificial intelligence will change the future of marketing, *Journal of the Academy of Marketing Science* 48, 24–42.
- Dietvorst, Berkeley J, Joseph P Simmons, and Cade Massey, 2015, Algorithm aversion: people erroneously avoid algorithms after seeing them err., *Journal of experimental psychology: General* 144, 114.
- Dietvorst, Berkeley J, Joseph P Simmons, and Cade Massey, 2018, Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them, *Management science* 64, 1155–1170.
- Dobbie, Will, Andres Liberman, Daniel Paravisini, and Vikram Pathania, 2021, Measuring bias in consumer lending, *The Review of Economic Studies* 88, 2799–2832.
- D’acunto, Francesco, Pulak Ghosh, and Alberto G Rossi, 2026, How costly are cultural biases? Evidence from FinTech, *Journal of Financial Economics* 175, 104202.
- Feldman, Michael, Sorelle A Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian, 2015, Certifying and removing disparate impact, in *proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, 259–268.
- Fügener, Andreas, Jörn Grahl, Alok Gupta, and Wolfgang Ketter, 2021, Will humans-in-the-loop become borgs? Merits and pitfalls of working with AI, *Management Information Systems Quarterly (MISQ)-Vol* 45.
- 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, 5–47.
- Fuster, Andreas, Matthew Plosser, Philipp Schnabl, and James Vickery, 2019, The role of technology in mortgage lending, *The review of financial studies* 32, 1854–1899.
- Gentzkow, Matthew, Bryan Kelly, and Matt Taddy, 2019, Text as data, *Journal of Economic Literature* 57, 535–574.
- Giannetti, Mariassunta, Mike Burkart, and Tore Ellingsen, 2011, What You Sell is What You Lend? Explaining Trade Credit Contracts, *The Review of Financial Studies* 24, 1261–1298.
- Goh, Kim Huat, Le Wang, Adrian Yong Kwang Yeow, Hermione Poh, Ke Li, Joannas Jie Lin Yeow, and Gamaliel Yu Heng Tan, 2021, Artificial intelligence in sepsis early prediction and diagnosis using unstructured data in healthcare, *Nature communications* 12, 711.
- Goodyear, Victoria A, Charlotte Kerner, and Mikael Quennerstedt, 2019, Young people’s uses of wearable healthy lifestyle technologies; surveillance, self-surveillance and resistance, *Sport, education and society* 24, 212–225.
- Grove, William M, David H Zald, Boyd S Lebow, Beth E Snitz, and Chad Nelson, 2000, Clinical versus mechanical prediction: a meta-analysis., *Psychological assessment* 12, 19.
- Gu, Shihao, Bryan Kelly, and Dacheng Xiu, 2020, Empirical asset pricing via machine learning, *The Review of Financial Studies* 33, 2223–2273.
- Guild, James, 2017, Fintech and the Future of Finance, *Asian Journal of Public Affairs* 17–20.
- Guo, Husheng, and Wenjian Wang, 2015, An active learning-based SVM multi-class classification model, *Pattern recognition* 48, 1577–1597.
- Hassan, Tarek A, Stephan Hollander, Laurence Van Lent, and Ahmed Tahoun, 2019, Firm-level political risk: Measurement and effects, *The quarterly journal of economics* 134, 2135–2202.
- Heider, Florian, and Roman Inderst, 2012, Loan prospecting, *The Review of Financial Studies* 25, 2381–2415.
- Hertzberg, Andrew, Jose Maria Liberti, and Daniel Paravisini, 2010, Information and incentives inside the firm: Evidence from loan officer rotation, *The Journal of Finance* 65, 795–828.
- Hirshleifer, David, Sonya Seongyeon Lim, and Siew Hong Teoh, 2009, Driven to distraction: Extraneous events and underreaction to earnings news, *The journal of finance* 64, 2289–2325.
- Hoffman, Mitchell, Lisa B Kahn, and Danielle Li, 2018, Discretion in hiring, *The Quarterly Journal of Economics* 133, 765–800.
- Holmström, Bengt, 1999, Managerial incentive problems: A dynamic perspective, *The review of Economic studies* 66, 169–182.
- Horton, John J, 2017, The effects of algorithmic labor market recommendations: Evidence from a field experiment, *Journal of Labor Economics* 35, 345–385.
- Hu, Xiyang, Yan Huang, Beibei Li, and Tian Lu, 2025, Human–algorithmic bias: Source, evolution, and impact, *Management Science* .

- Huang, Ruidi, Erik J Mayer, and Darius P Miller, 2024, Gender bias in promotions: Evidence from financial institutions, *The Review of Financial Studies* 37, 1685–1728.
- Ibrahim, Rouba, Song-Hee Kim, and Jordan Tong, 2021, Eliciting human judgment for prediction algorithms, *Management Science* 67, 2314–2325.
- Jacovi, Alon, Ana Marasović, Tim Miller, and Yoav Goldberg, 2021, Formalizing trust in artificial intelligence: Prerequisites, causes and goals of human trust in AI, in *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 624–635.
- Jagtiani, Julapa, and Catharine Lemieux, 2019, The roles of alternative data and machine learning in fintech lending: Evidence from the LendingClub consumer platform, *Financial Management* 48, 1009–1029.
- Jansen, Mark, Hieu Quang Nguyen, and Amin Shams, 2025, Rise of the machines: The impact of automated underwriting, *Management Science* 71, 955–975.
- Klare, Brendan F, Mark J Burge, Joshua C Klontz, Richard W Vorder Bruegge, and Anil K Jain, 2012, Face recognition performance: Role of demographic information, *IEEE Transactions on information forensics and security* 7, 1789–1801.
- Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer, 2015, Prediction policy problems, *American Economic Review* 105, 491–495.
- Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Ashesh Rambachan, 2018, Algorithmic fairness, *AEA Papers and Proceedings* 108, 22–27.
- Krakowski, Sebastian, Johannes Luger, and Sebastian Raisch, 2023, Artificial intelligence and the changing sources of competitive advantage, *Strategic Management Journal* 44, 1425–1452.
- Kumar, Pawan, and Manmohan Sharma, 2022, Data, machine learning, and human domain experts: none is better than their collaboration, *International Journal of Human–Computer Interaction* 38, 1307–1320.
- Lebovitz, Sarah, Natalia Levina, and Hila Lifshitz-Assaf, 2021, Is AI ground truth really true? The dangers of training and evaluating AI tools based on experts’ know-what, *MIS quarterly* 45, 1501–1526.
- Li, Chunxiao, Hongchang Wang, Songtao Jiang, and Bin Gu, 2024, The effect of ai-enabled credit scoring on financial inclusion: Evidence from an underserved population of over one million, *MIS Quarterly* 48, 1803–1834.
- Liang, Annie, Jay Lu, Xiaosheng Mu, and Kyohei Okumura, 2026, Algorithm design: A fairness-accuracy frontier, *Journal of Political Economy* .
- Lin, Mingfeng, and Siva Viswanathan, 2016, Home bias in online investments: An empirical study of an online crowdfunding market, *Management science* 62, 1393–1414.
- Liu, Meng, Xiaocheng Tang, Siyuan Xia, Shuo Zhang, Yuting Zhu, and Qianying Meng, 2026, Algorithm aversion: Evidence from ridesharing drivers, *Management Science* 72, 193–203.
- Lou, Bowen, and Lynn Wu, 2021, AI on drugs: Can artificial intelligence accelerate drug development? Evidence from a large-scale examination of bio-pharma firms, *MIS quarterly* 45, 1451–1482.
- Loufield, Ethan, Dennis Ferenzy, and Tess Johnson, 2018, Accelerating financial inclusion with new data, *Center for Financial Inclusion* 6.
- Lu, Tian, and Yingjie Zhang, 2025, 1+ 1 > 2? Information, humans, and machines, *Information Systems Research* 36, 394–418.
- Lu, Tian, Yingjie Zhang, and Beibei Li, 2023, Profit vs. equality? The case of financial risk assessment and a new perspective on alternative data, *MIS quarterly* 47, 1517–1556.
- Luo, Xueming, Marco Shaojun Qin, Zheng Fang, and Zhe Qu, 2021, Artificial intelligence coaches for sales agents: Caveats and solutions, *Journal of Marketing* 85, 14–32.
- March, James G, 1991, Exploration and exploitation in organizational learning, *Organization science* 2, 71–87.
- Obermeyer, Ziad, Brian Powers, Christine Vogeli, and Sendhil Mullainathan, 2019, Dissecting racial bias in an algorithm used to manage the health of populations, *Science* 366, 447–453.
- Petersen, Mitchell A., and Raghuram G. Rajan, 1994, The Benefits of Lending Relationships: Evidence from Small Business Data, *The Journal of Finance* 49, 3–37.
- Phelps, Edmund S, 1972, The statistical theory of racism and sexism, *The american economic review* 62, 659–661.
- Sawyer, Jack, 1966, Measurement and prediction, clinical and statistical., *Psychological bulletin* 66, 178.
- Schwab, Stewart, 1986, Is statistical discrimination efficient?, *The American Economic Review* 76, 228–234.

- Schwartz, Robert A., 1974, An Economic Model of Trade Credit, *Journal of Financial and Quantitative Analysis* 9, 643–657.
- Siau, Keng, and Weiyu Wang, 2018, Building trust in artificial intelligence, machine learning, and robotics, *Cutter business technology journal* 31, 47.
- Sun, Jiankun, Dennis J Zhang, Haoyuan Hu, and Jan A Van Mieghem, 2022, Predicting human discretion to adjust algorithmic prescription: A large-scale field experiment in warehouse operations, *Management Science* 68, 846–865.
- Sweeney, Latanya, 2013, Discrimination in online ad delivery, *Communications of the ACM* 56, 44–54.
- Teodorescu, Mike HM, Lily Morse, Yazeed Awwad, and Gerald C Kane, 2021, Failures of fairness in automation require a deeper understanding of human–ml augmentation, *MIS quarterly* 45, 1483–1500.
- Tetlock, Paul C, Maytal Saar-Tsechansky, and Sofus Macskassy, 2008, More than words: Quantifying language to measure firms’ fundamentals, *The journal of finance* 63, 1437–1467.
- Te’eni, Dov, Inbal Yahav, Alexely Zagalsky, David Schwartz, Gahl Silverman, Daniel Cohen, Yossi Mann, and Dafna Lewinsky, 2026, Reciprocal human-machine learning: A theory and an instantiation for the case of message classification, *Management Science* 72, 167–192.
- Tong, Siliang, Nan Jia, Xueming Luo, and Zheng Fang, 2021, The Janus face of artificial intelligence feedback: Deployment versus disclosure effects on employee performance, *Strategic Management Journal* 42, 1600–1631.
- Udell, Gregory F, 1989, Loan quality, commercial loan review and loan officer contracting, *Journal of Banking & Finance* 13, 367–382.
- Vallee, Boris, and Yao Zeng, 2019, Marketplace lending: A new banking paradigm?, *The Review of Financial Studies* 32, 1939–1982.
- von Walter, Benjamin, Dietmar Kremmel, and Bruno Jäger, 2022, The impact of lay beliefs about AI on adoption of algorithmic advice, *Marketing Letters* 33, 143–155.
- Wang, Hongchang, Yingjie Zhang, and Tian Lu, 2025, The Power of Disagreement: A Field Experiment to Investigate Human–Algorithm Collaboration in Loan Evaluations, *Management Science* .
- Wang, Wen, Mochen Yang, and Tianshu Sun, 2023, Human-AI co-creation in product ideation: The dual view of quality and diversity, *Available at SSRN 4668241* .
- Zhang, Mengxia, Tianshu Sun, Lan Luo, and Joseph Golden, 2021, Consumer and AI Co-creation: when and why human participation improves AI creation, *USC Marshall School of Business Research Paper Sponsored by iORB, No. Forthcoming* .
- Zhao, Eric Yanfei, and Tyler Wry, 2016, Not all inequality is equal: Deconstructing the societal logic of patriarchy to understand microfinance lending to women, *Academy of management Journal* 59, 1994–2020.

Figures and Tables

Figure 1. Sample Credit Report Generated by the FinTech Scoring System

This figure presents a sample credit report generated by the FinTech scoring system. The left section displays basic customer information, including customer ID, name, phone number, and national ID number. The right section shows the machine-generated credit score and the corresponding recommendation (approval or rejection). The lower section provides detailed risk alerts, listing potential sources of credit risk identified by the algorithm.



Figure 2. The Loan Approval Process of the Auto Finance Leasing Company

This figure illustrates the credit approval process of the auto finance leasing company following the adoption of the FinTech scoring system. The process begins with the applicant selecting a vehicle and submitting a loan application. A front-line credit officer conducts a preliminary review to verify the authenticity of submitted information. The FinTech system then generates a credit score and an approval recommendation. A senior credit officer performs the final review, referencing the machine-generated score and, if necessary, consulting supplementary materials such as field investigation reports and telephone interview records. The final loan decision is made by the senior officer. Red boxes highlight the stages where decisions are made by the algorithm and by the human officer, which are the focus of our empirical analysis.

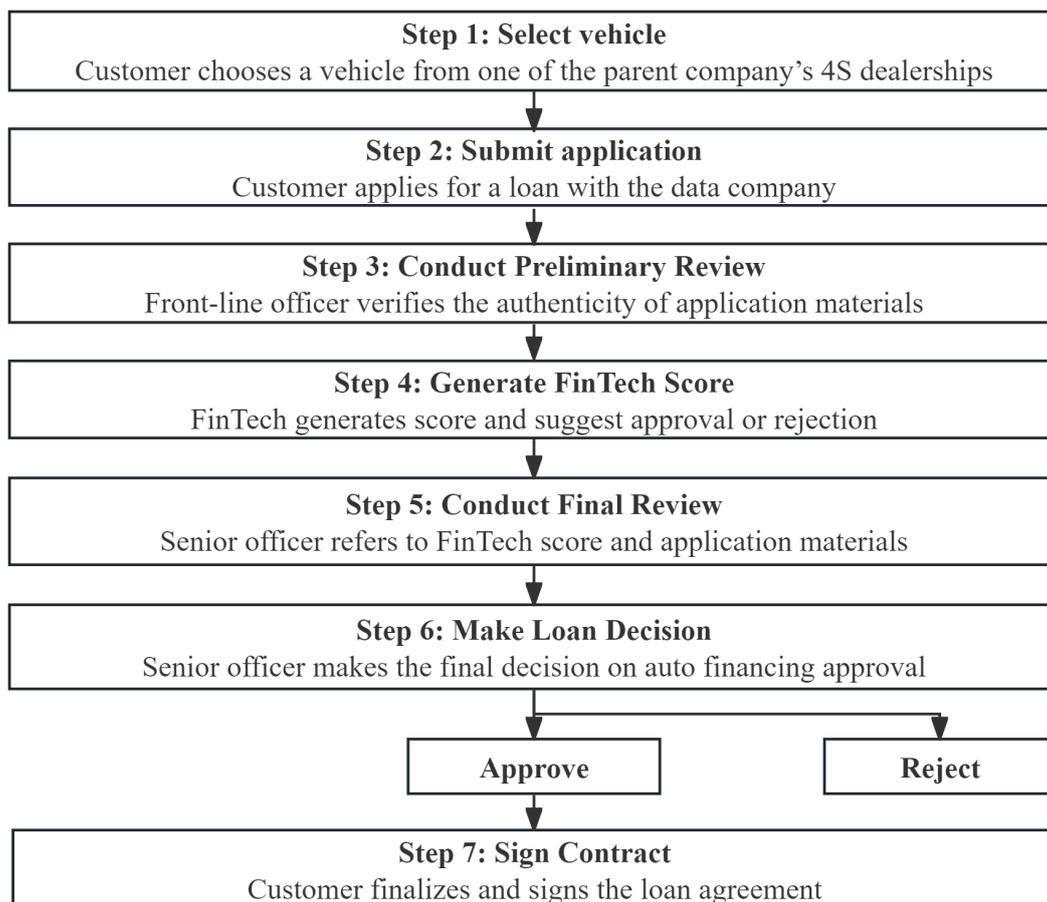


Figure 3. Distribution of FinTech Credit Scores

This figure illustrates the distribution of FinTech credit scores for all passenger vehicle applicants after the implementation of the FinTech scoring system in July 2020, including both approved and rejected applications.

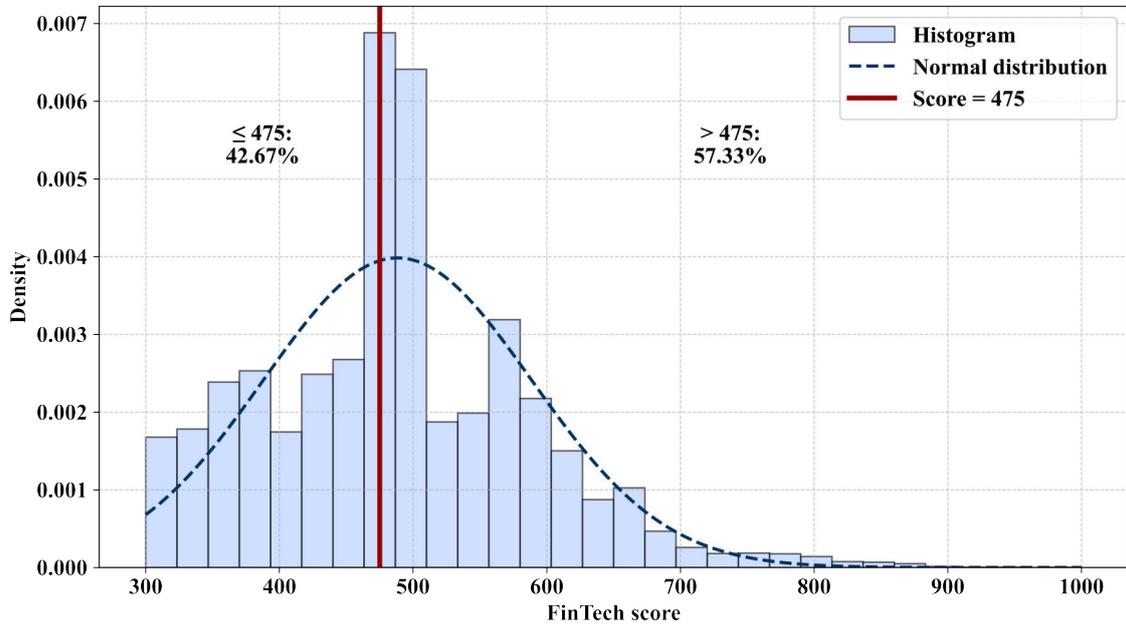


Figure 4. Approval Rate by FinTech Credit Score Bin

This figure shows the loan approval rate within each FinTech credit score bin for passenger vehicle applicants after its implementation in July 2020. For each score range, it displays the proportion of approved applications relative to the total number of applications in that bin.

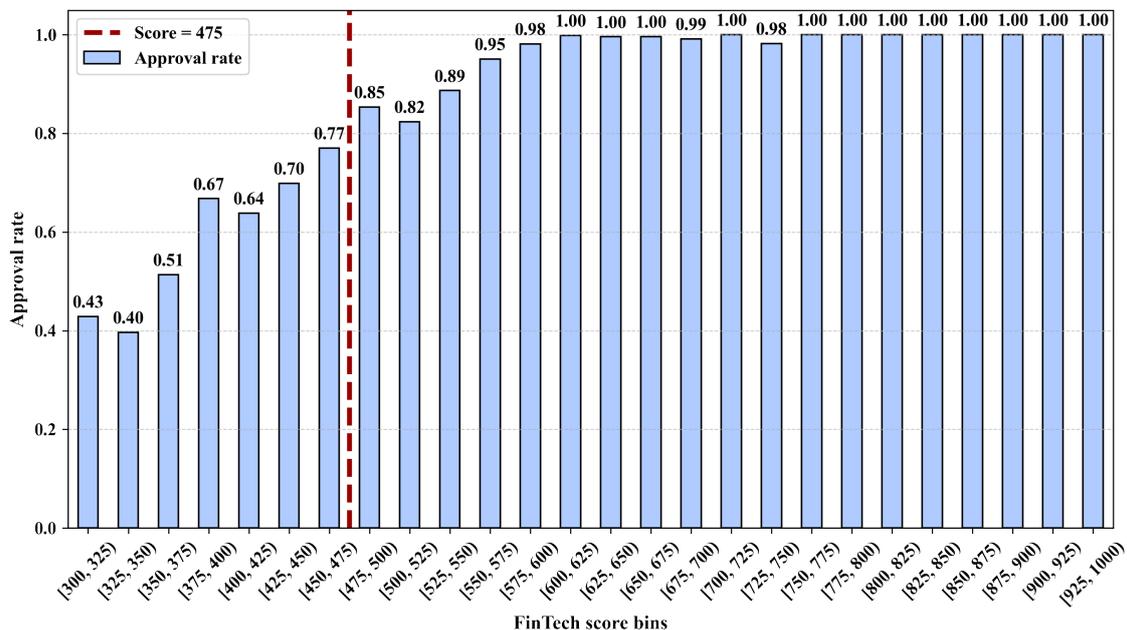


Figure 5. Approval Rates by FinTech Credit Score: Early vs. Late Post-Adoption Period

This figure compares loan approval rates across FinTech credit score bins between two post-adoption periods: the early period and the late period. Approval rates are computed within score bins and smoothed using local averaging to highlight underlying patterns. The vertical red dashed line indicates the score cutoff at 475.

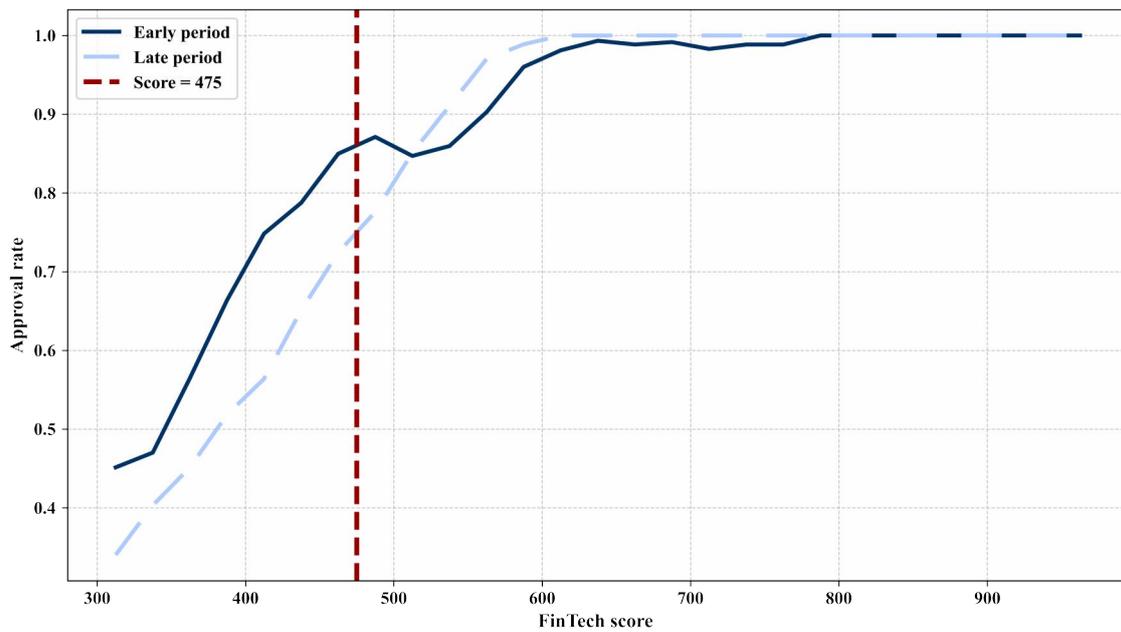
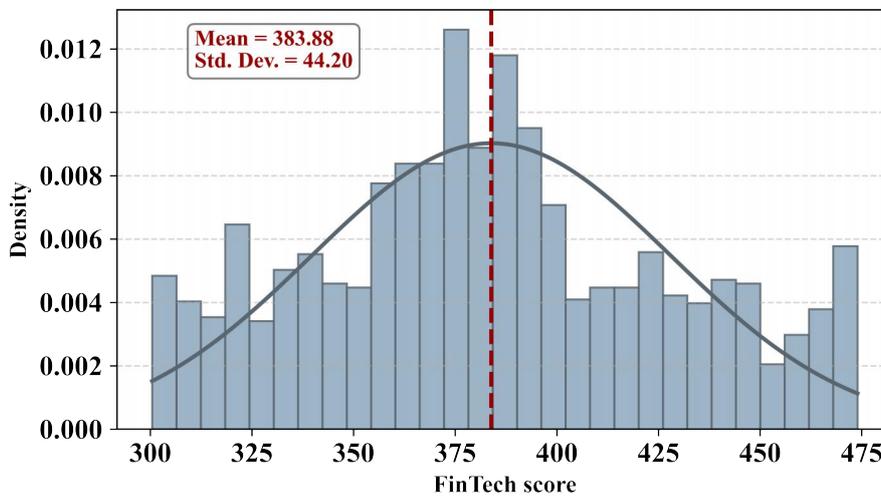


Figure 6. Distribution of FinTech Credit Scores by Post-Adoption Period

This figure presents the distribution of FinTech credit scores for passenger vehicle applicants whose loan applications were approved after the adoption of the FinTech scoring system. The post-adoption period is divided into two equal-length intervals, referred to as the early period in subfigure (a) and the late period in subfigure (b). Each subplot displays the score distribution for approved loans in the corresponding period. The vertical red line indicates the cutoff score of 475.

(a) Post-adoption: Early period



(b) Post-adoption: Late period

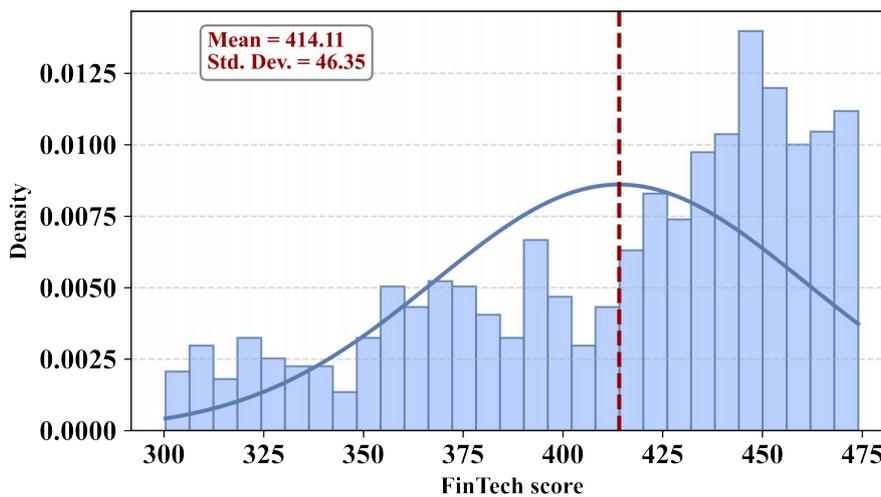


Figure 7. Average Audit Time Before and After FinTech Adoption

This figure compares the average audit time (in minutes) across different stages and subsamples surrounding the adoption of the FinTech credit scoring system, using the full sample of passenger vehicle loan applications, including both approved and rejected cases. Panel A contrasts audit time before and after FinTech adoption. Panel B focuses on the post-adoption period and compares applicants with credit scores below and above the cutoff of 475. Panel C further splits the post-adoption sample into two equal-length intervals, referred to as the early period and the late period. Panel D jointly considers time and score heterogeneity by comparing audit time across early versus late periods and low-score versus high-score applicants in the post-adoption period. Each bar reports the mean audit time, with the corresponding standard deviation displayed above the bar.

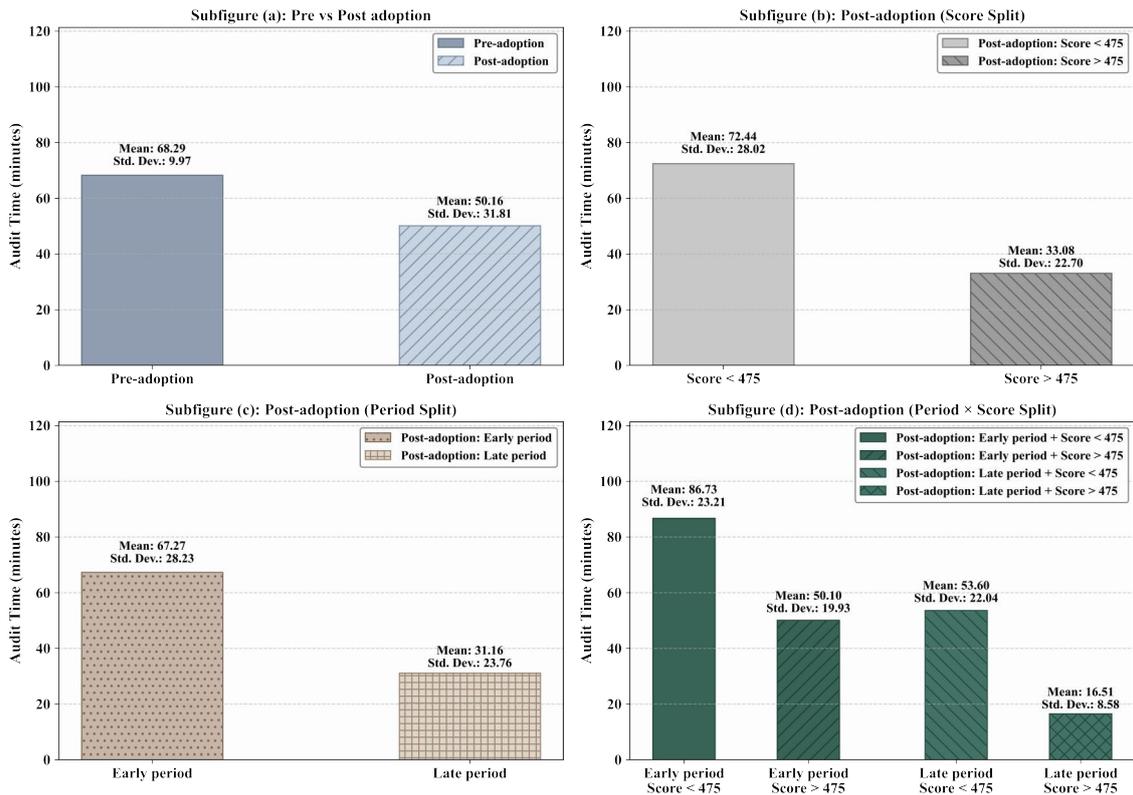


Figure 8. Average Audit Time across FinTech Credit Score Bins: Early vs. Late Post-Adoption Period

This figure compares average audit time (in minutes) across FinTech credit score bins between the early and late periods following the adoption of the FinTech credit scoring system, using the full sample of passenger vehicle loan applications, including both approved and rejected cases. Audit times are first averaged within score bins and then smoothed to highlight underlying patterns along the score distribution. The vertical red dashed line indicates the score cutoff at 475.

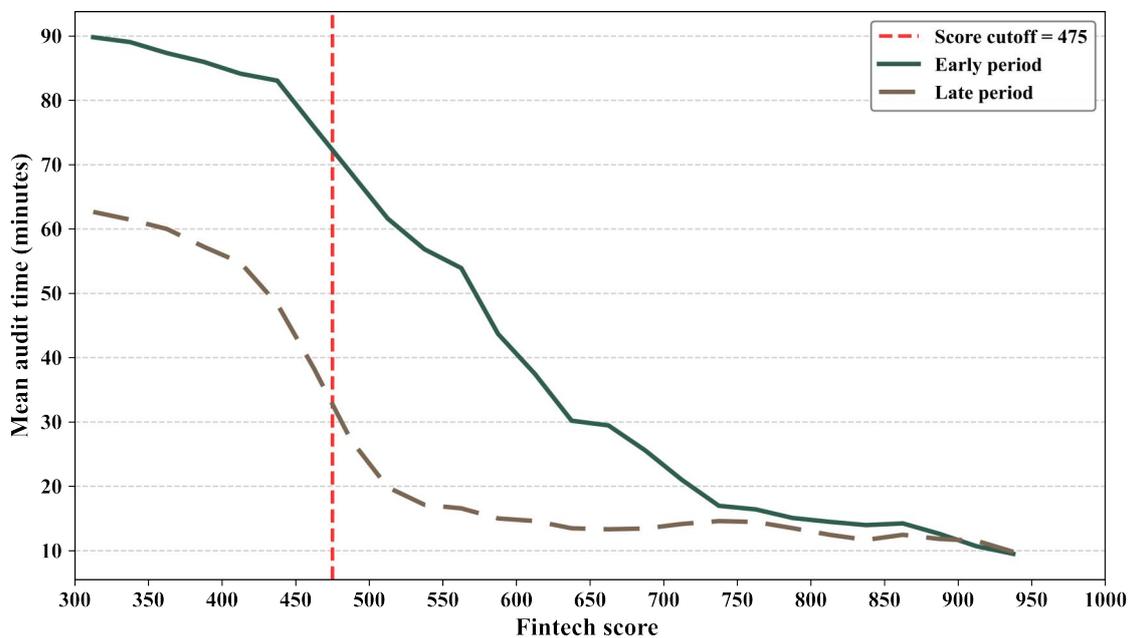


Figure 9. Distribution of Rescued Orders by Promotion Status and Post-Adoption Period

This figure presents the distribution of FinTech credit scores for rescued orders in the full sample of passenger vehicle loan applications, including both approved and rejected cases, stratified by promotion status and time period following the adoption of the FinTech scoring system. Subfigure (a) compares the score distributions of rescued orders between the promotion period and the non-promotion period. Subfigure (b) further splits the sample by post-adoption timing, distinguishing between the early and late periods within promotion and non-promotion regimes. In each subplot, the vertical red line denotes the sample mean of FinTech credit scores.

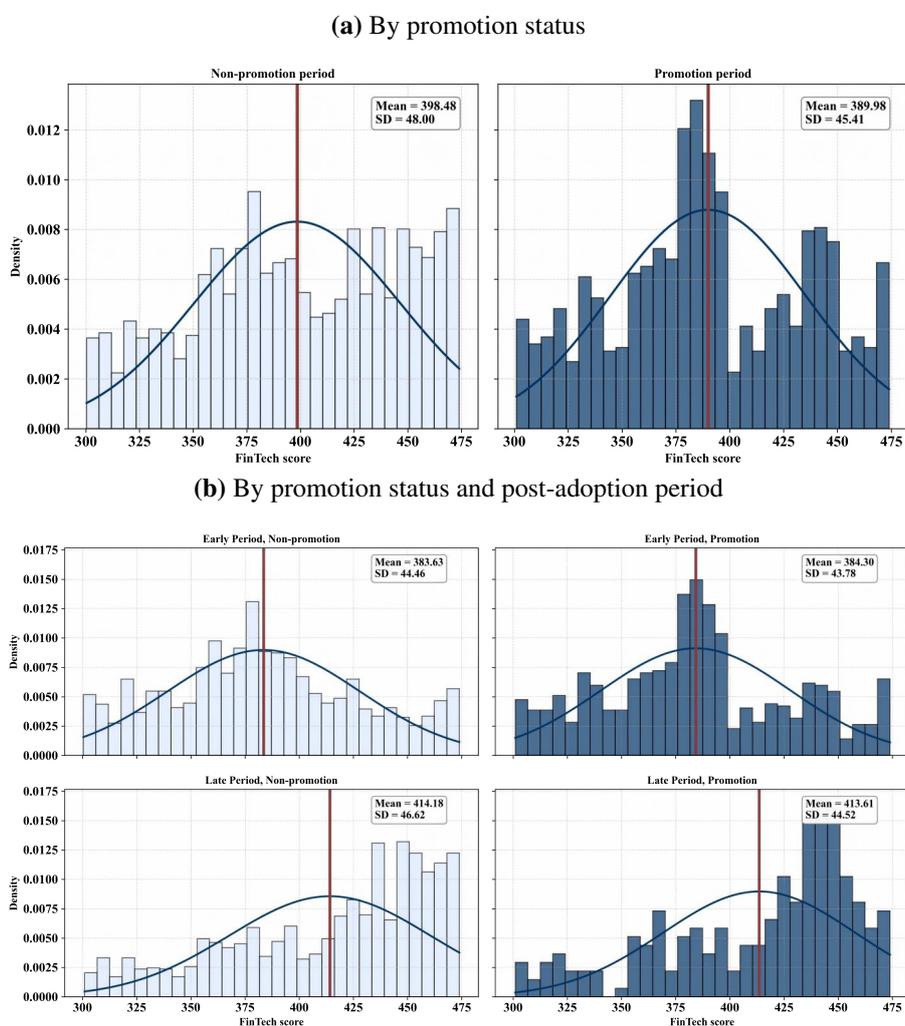
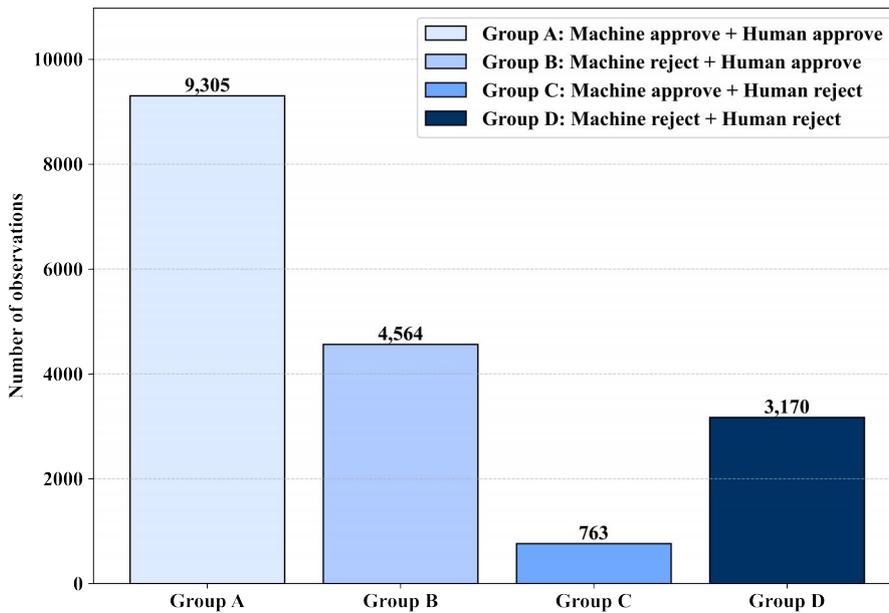


Figure 10. Sample Composition after FinTech Adoption

This figure reports the distribution of loan applications across four decision outcome groups in the full sample of passenger vehicle loan applications following the adoption of the FinTech credit scoring system. Subfigure (a) presents the total number of observations in each outcome group in the post-adoption sample. The four groups are defined by the joint decisions of the machine and the human reviewer: machine approve and human approve (Group A), machine reject and human approve (Group B), machine approve and human reject (Group C), and machine reject and human reject (Group D). Subfigure (b) further decomposes the post-adoption sample by time, distinguishing between the early period and the late period after adoption, and reports the corresponding group-level counts in each period.

(a) Sample size by group



(b) Sample size by group and period

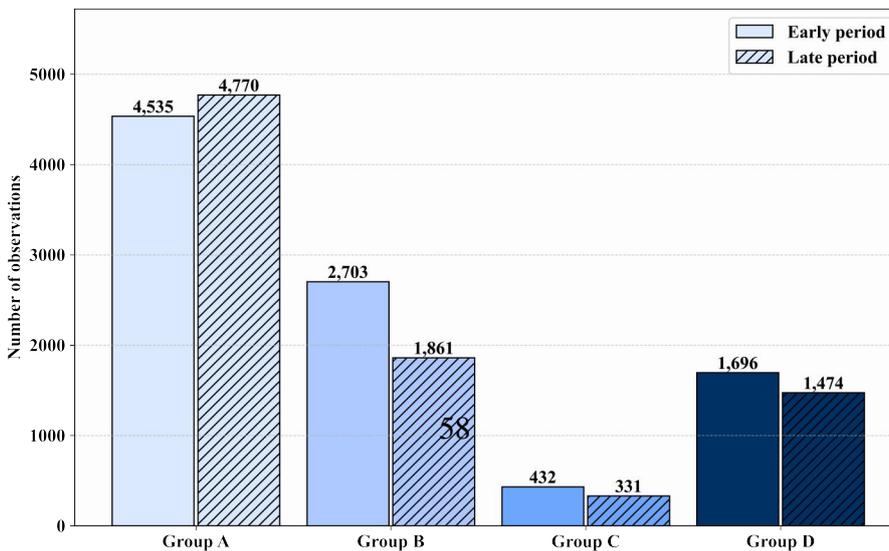


Table 1. Effects of FinTech Credit Scoring Adoption on Default Probability and Loan Profit

This table examines the effects of adopting the FinTech credit scoring system on loan outcomes. The dependent variables include a default indicator, which equals one if a loan defaults and zero otherwise, and the loan profit ratio, defined as profit divided by initial investment. The treatment variable, *Treat*, equals one for loan applications submitted after the implementation of the FinTech scoring system—March 2020 for commercial vehicles and July 2020 for passenger vehicles—and zero otherwise. Control variables include contract-level, borrower-level, and vehicle-level characteristics. “Uncensored” loans refer to loans whose original maturities do not extend beyond the end of the sample period, while “completed” loans are those that are fully completed within the sample period. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. Coefficients marked with *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle Variable Sample	Passenger + Commercial					
	Default			Loan Profit Ratio		
	All	Uncensored	Completed	All	Uncensored	Completed
	(1)	(2)	(3)	(4)	(5)	(6)
Treat	-0.165*** (-6.87)	-0.165*** (-6.89)	-0.142*** (-4.86)	0.037*** (4.27)	0.034*** (4.07)	0.026*** (3.70)
Log(Contract Duration)	0.190*** (5.48)	0.264*** (11.12)	0.722*** (5.64)	-0.012 (-0.65)	-0.097*** (-5.91)	-0.243*** (-5.08)
Down Payment Ratio	0.000 (0.50)	0.000 (1.06)	0.000 (0.44)	0.002*** (12.92)	0.002*** (9.87)	0.002*** (9.18)
Log(Contract Financing)	0.085*** (9.56)	0.074*** (4.85)	0.078*** (6.40)	-0.064*** (-9.05)	-0.055*** (-4.69)	-0.066*** (-8.43)
Contract Interest Rate	0.013*** (11.64)	0.010*** (10.72)	0.009*** (14.49)	0.002*** (3.58)	0.005*** (7.36)	0.003*** (4.02)
New Vehicle	0.014 (0.91)	-0.003 (-0.23)	0.096*** (3.28)	-0.006 (-0.67)	0.004 (0.51)	-0.025 (-1.37)
Business Use	0.139*** (2.00)	0.059 (0.57)	0.150 (1.58)	-0.036 (-0.71)	0.025 (0.33)	0.002 (0.02)
Vehicle Brand	-0.023*** (-3.62)	-0.030*** (-3.74)	-0.035*** (-3.97)	0.014*** (3.63)	0.020*** (4.45)	0.024*** (5.60)
Log(Client Wage)	-0.000 (-0.04)	0.010** (2.51)	0.008 (1.62)	0.001 (0.17)	-0.008** (-2.11)	-0.004 (-1.00)
Client Workplace	-0.104*** (-10.57)	-0.075*** (-5.37)	-0.091*** (-7.12)	0.028*** (9.23)	0.023*** (4.95)	0.026*** (5.79)
Client House Ownership	-0.194*** (-12.47)	-0.196*** (-10.44)	-0.207*** (-12.77)	0.103*** (9.12)	0.094*** (6.31)	0.110*** (9.40)
Married Client	-0.033*** (-7.55)	-0.009 (-1.48)	-0.019*** (-3.66)	0.014*** (6.03)	0.006* (1.79)	0.010*** (3.64)
Has Children	0.009*** (2.78)	0.002 (0.46)	0.003 (0.90)	-0.008*** (-5.74)	-0.006*** (-2.97)	-0.007*** (-4.25)
Local Client	-0.113*** (-8.06)	-0.206*** (-6.85)	-0.242*** (-8.72)	0.023*** (3.07)	0.068*** (4.99)	0.047*** (3.22)
Contract Province × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Officer × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Agent × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	134,467	99,013	105,492	134,467	99,013	105,492
R ²	0.244	0.305	0.467	0.170	0.171	0.263

Table 2. Effects of FinTech Credit Scoring Adoption on Coarse-Group Information Reliance

This table examines the effects of adopting the FinTech credit scoring system on coarse-group information reliance. The dependent variables include Coarse-Group Information Reliance (CGIR), log(client wage), client workplace, client house ownership, married client, has children, and local client. The treatment variable, *Treat*, equals one for loan applications submitted after the implementation of the FinTech scoring system—March 2020 for commercial vehicles and July 2020 for passenger vehicles—and zero otherwise. Control variables include contract-level and vehicle-level characteristics. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. Coefficients marked with *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Variable	Passenger + Commercial						
	CGIR	Log(Client Wage)	Client Workplace	Client House Ownership	Married Client	Has Children	Local Client
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treat	0.330*** (12.62)	0.279*** (3.69)	0.279*** (7.90)	0.104*** (6.29)	0.041* (1.95)	0.149*** (5.62)	-0.040*** (-3.00)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Contract Province × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Officer × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Agent × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	134,467	134,467	134,467	134,467	134,467	134,467	134,467
R^2	0.146	0.591	0.155	0.451	0.038	0.072	0.567

Table 3. The Performance of Human–Machine Collaboration on the Rescued Group

This table investigates the effect of human–machine collaboration in credit approval exceptions using a subsample of passenger vehicle loans restricted to the interval $[475 - \text{MBW}, 475 + \text{MBW}]$. MBW denotes the margin bandwidth used to construct a locally comparable sample around the approval cutoff. The dependent variables include a default indicator (equal to one if a loan defaults and zero otherwise), the loan profit ratio (defined as profit divided by initial investment), Coarse-Group Information Reliance (CGIR), $\log(\text{client wage})$, client workplace, client house ownership, married client, has children, and local client. *Rescued* is an indicator variable equal to one for applications initially rejected by the machine but subsequently approved by a human, and zero otherwise. Control variables include contract- and vehicle-level characteristics; for default-related outcomes, borrower-level controls are additionally included. Standard errors are two-way clustered at the officer-by-year level, with t -statistics reported in parentheses. Coefficients marked with *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Sample	Comparable Passenger Vehicle					
Panel A: Loan outcomes						
Variables	Default	Loan Profit Ratio			CGIR	
	(1)	(2)	(2)	(2)	(3)	(3)
Rescued	0.000 (0.03)	-0.003 (-0.65)			-0.038* (-1.80)	
Observations	4,961	4,961			4,961	
R^2	0.095	0.118			0.291	
Panel B: Borrower Characteristics						
Variables	Log(Client Wage)	Client Workplace	Client House Ownership	Married Client	Has Children	Local Client
	(1)	(2)	(3)	(4)	(5)	(6)
Rescued	-0.100*** (-2.93)	-0.019 (-1.25)	-0.040*** (-3.41)	-0.057*** (-6.31)	-0.001 (-0.12)	0.085*** (4.84)
Observations	4,961	4,961	4,961	4,961	4,961	4,961
R^2	0.345	0.163	0.258	0.046	0.068	0.338
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Contract Province \times Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Officer \times Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Agent \times Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 4. Dynamic Effects of Human–Machine Collaboration for Rescued Loans

This table examines the dynamic effects of human–machine collaboration in credit approval exceptions, using a subsample of passenger vehicle loans restricted to the interval $[475 - \text{MBW}, 475 + \text{MBW}]$. MBW denotes the margin bandwidth used to construct a locally comparable sample around the approval cutoff. The dependent variables include a default indicator (equal to one if a loan defaults and zero otherwise), the loan profit ratio (defined as profit divided by initial investment), Coarse-Group Information Reliance (CGIR), $\log(\text{client wage})$, client workplace, client house ownership, married client, has children, and local client. *Late period*, is an indicator equal to one for loans originated in the late post-adoption period and zero for those originated in the early post-adoption period. *Rescued* is an indicator variable equal to one for applications initially rejected by the machine but subsequently approved by a human, and zero otherwise. Control variables include contract- and vehicle-level characteristics; for default-related outcomes, borrower-level controls are additionally included. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. Coefficients marked with *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Sample	Comparable Passenger Vehicle					
<i>Panel A: Loan outcomes</i>						
Variables	Default	Loan Profit Ratio			CGIR	
	(1)	(2)			(3)	
Rescued \times Late period	-0.054*** (-2.76)	0.004 (0.47)			0.120*** (3.70)	
Rescued	0.082** (2.20)	-0.009 (-0.63)			-0.225*** (-4.47)	
Late period	0.027** (2.16)	-0.002 (-0.32)			-0.037** (-2.14)	
Observations	4,961	4,961			4,961	
R^2	0.097	0.118			0.293	
<i>Panel B: Borrower Characteristics</i>						
Variables	Log(Client Wage)	Client Workplace	Client House Ownership	Married Client	Has Children	Local Client
	(1)	(2)	(3)	(4)	(5)	(6)
Rescued \times Late period	0.247*** (4.74)	-0.018 (-0.95)	0.098*** (2.76)	0.035 (1.60)	-0.001 (-0.03)	0.031 (0.78)
Rescued	-0.455*** (-5.92)	0.008 (0.34)	-0.203*** (-3.88)	-0.114*** (-3.28)	-0.004 (-0.12)	0.039 (0.59)
Late period	-0.204*** (-3.86)	0.011 (1.00)	0.019 (1.24)	0.003 (0.15)	0.017 (1.41)	-0.016 (-0.60)
Observations	4,961	4,961	4,961	4,961	4,961	4,961
R^2	0.352	0.164	0.262	0.047	0.069	0.337
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Contract Province \times Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Officer \times Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Agent \times Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 5. Effects of FinTech Credit Scoring Adoption on Approval Text Characteristics

This table examines how the adoption of the FinTech credit scoring system affects the textual characteristics of loan approval opinions for passenger vehicle loans. The sample is restricted to approved applications only. The dependent variables capture different dimensions of approval text content: the logarithm of text length, information density, a dummy indicator for the presence of fuzzy or vague expressions, and the average semantic similarity across approval opinions. *Post* is an indicator equal to one for applications submitted after the introduction of the FinTech scoring system. *Late Period* denotes the later stage of the post-adoption period. Control variables include contract-, borrower-level and vehicle-level characteristics. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. Coefficients marked with *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle	Passenger			
Sample	Approved Applications			
Variables	Log(Text Length)	Information Density	Fuzzy Dummy	Textual Similarity
<i>Panel A: Average Effects of FinTech Adoption</i>				
Period	Full			
	(1)	(2)	(3)	(4)
Post	−0.011* (−1.75)	−0.001 (−0.87)	0.001*** (3.42)	0.002 (0.31)
Observations	19,209	19,209	19,209	19,209
R^2	0.010	0.009	0.019	0.125
<i>Panel B: Dynamic Effects in the Late Post-Adoption Period</i>				
Period	Post-Adoption			
Late Period	−0.024*** (−4.90)	−0.001*** (−3.02)	0.002*** (9.66)	0.015*** (2.70)
Observations	13,869	13,869	13,869	13,869
R^2	0.016	0.011	0.032	0.088
Controls	Yes	Yes	Yes	Yes
Contract Province × Year FE	Yes	Yes	Yes	Yes
Officer × Year FE	Yes	Yes	Yes	Yes
Agent × Year FE	Yes	Yes	Yes	Yes

Table 6. Dynamic Changes in Approval Text Characteristics near the FinTech Score Cutoff

This table examines the dynamic evolution of approval text characteristics following the adoption of the FinTech credit scoring system for passenger vehicle loans. The sample is restricted to approved applications only. Textual outcomes include the logarithm of text length, information density, a dummy indicator for the presence of fuzzy or vague expressions, and the average semantic similarity across approval opinions. *Late period*, is an indicator equal to one for loans originated in the late post-adoption period and zero for those originated in the early post-adoption period. *Near cutoff*³⁰ indicates whether the FinTech score falls within [445,475]. *Log score distance* is the logarithm of the absolute distance between the score and the cutoff. All specifications include contract-, borrower-, and vehicle-level controls, as well as contract-province-by-year, officer-by-year, and agent-by-year fixed effects. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle	Passenger			
Sample	Approved			
Variables	Log(Text Length)	Info. Density	Fuzzy Dummy	Textual Similarity
Panel A: Near Cutoff³⁰				
Late Period × Near cutoff ³⁰	−0.173*** (−9.31)	−0.003** (−2.26)	0.016*** (15.82)	0.016 (0.75)
Late Period	−0.006 (−0.98)	−0.000 (−0.21)	−0.000 (−1.25)	0.021 (1.39)
Near cutoff ³⁰	0.169*** (6.15)	0.003* (1.88)	−0.016*** (−15.20)	0.005 (0.15)
Observations	4,543	4,543	4,543	4,543
R ²	0.106	0.025	0.294	0.191
Variables	Log(Text Length)	Info. Density	Fuzzy Dummy	Textual Similarity
Panel B: Log(Score Distance)				
Late Period × Log(Score Distance)	0.065*** (7.33)	0.001* (1.76)	−0.006*** (−14.59)	−0.013 (−1.12)
Late Period	−0.306*** (−7.83)	−0.005* (−1.88)	0.027*** (14.38)	0.078* (1.72)
Log(Score Distance)	−0.062*** (−4.88)	−0.001 (−1.12)	0.006*** (14.95)	0.008 (0.42)
Observations	4,543	4,543	4,543	4,543
R ²	0.084	0.025	0.216	0.193
Controls	Yes	Yes	Yes	Yes
Contract Province × Year FE	Yes	Yes	Yes	Yes
Officer × Year FE	Yes	Yes	Yes	Yes
Agent × Year FE	Yes	Yes	Yes	Yes

Table 7. Effects of FinTech Credit Scoring Adoption on Audit Time

This table examines the effects of the adoption of the FinTech credit scoring system on audit time for passenger vehicle loan applications. The sample includes both approved and rejected applications. The dependent variable is the logarithm of audit time measured in minutes. *Post* is an indicator equal to one for applications submitted after the introduction of the FinTech scoring system. *Score > 475* indicates whether an application's FinTech credit score exceeds the approval cutoff of 475. *Late Period* denotes the later half of the post-adoption period. *Near cutoff*¹⁰, *Near cutoff*²⁰, and *Near cutoff*³⁰, are indicator variables equal to one if a loan's FinTech score falls within increasingly narrow intervals below the cutoff of 475, specifically [465,475], [455,475], and [445,475], respectively. We additionally consider the logarithm of the distance between the FinTech score and the cutoff, *Log(Score distance)*, as an alternative continuous measure of proximity to the cutoff. Control variables include contract-, borrower-level and vehicle-level characteristics. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. Coefficients marked with *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle	Passenger			
Sample	Both approved and rejected			
Panel A: Average Effects of FinTech Adoption on Audit Time				
Variables	Log(Audit Time)			
Period	Full	Post-Adoption		
	(1)	(2)	(3)	(4)
Post	-0.025* (-1.77)			
Score > 475		-0.734*** (-13.65)		-0.027 (-0.48)
Late Period			-0.890*** (-33.86)	-0.548*** (-20.62)
Late Period × Score > 475				-0.487*** (-14.36)
Observations	22,119	16,768	16,768	16,768
R ²	0.570	0.683	0.565	0.745
Panel B: Post-Adoption Cutoff Proximity				
Variables	Log(Audit Time)			
Period	Post-Adoption			
Near Cutoff measure	<i>NearCutoff</i> ¹⁰	<i>NearCutoff</i> ²⁰	<i>NearCutoff</i> ³⁰	Log(Score Distance)
Late Period × Near Cutoff	-0.308*** (-8.27)	-0.332*** (-12.53)	-0.267*** (-14.42)	0.138*** (15.37)
Late Period	-0.517*** (-22.55)	-0.477*** (-19.08)	-0.454*** (-16.32)	-1.053*** (-26.76)
Near Cutoff	0.310*** (7.59)	0.329*** (9.79)	0.251*** (9.84)	-0.130*** (-11.56)
Observations	6,700	6,700	6,700	6,700
R ²	0.510	0.531	0.533	0.536
Controls	Yes	Yes	Yes	Yes
Contract Province × Year FE	Yes	Yes	Yes	Yes
Officer × Year FE	Yes	Yes	Yes	Yes
Agent × Year FE	Yes	65	Yes	Yes

Table 8. Promotion Periods and the Distribution of Rescued Loans around the FinTech Score Cutoff

This table examines how promotion periods affect the distribution of rescued passenger vehicle loan applications around the FinTech score cutoff. The sample is restricted to passenger vehicle loans that were initially rejected by the machine but subsequently approved by a human reviewer. The dependent variables measure whether a rescued loan's FinTech score lies close to the approval cutoff of 475, using indicator variables for increasingly wider intervals below the cutoff: *Near Cutoff*¹⁰ ([465, 475]), *Near Cutoff*²⁰ ([455, 475]), and *Near Cutoff*³⁰ ([445, 475]), as well as the logarithm of the score distance from the cutoff. *Promotion* is an indicator for promotion periods, *Late Period* identifies the later half of the post-adoption period, and their interaction captures dynamic effects. All specifications include contract-, borrower-, and vehicle-level controls, as well as contract-province-by-year, officer-by-year, and agent-by-year fixed effects. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle Sample Variables	Passenger			
	Rescued			
	Near Cutoff ¹⁰	Near Cutoff ²⁰	Near Cutoff ³⁰	Log Score Distance
Panel A: Baseline Effects of Promotion Periods				
Promotion	-0.038*** (-3.01)	-0.069*** (-5.35)	-0.060*** (-2.79)	0.234*** (4.14)
Observations	4,543	4,543	4,543	4,543
R ²	0.046	0.058	0.056	0.051
Panel B: Dynamic Effects of Promotion Periods				
Late Period × Promotion	-0.052*** (-3.19)	-0.068*** (-2.93)	-0.072* (-1.90)	0.121 (1.56)
Promotion	0.045** (2.39)	0.048 (1.55)	0.085** (2.03)	-0.052 (-0.63)
Late Period	0.066*** (7.14)	0.132*** (9.42)	0.238*** (15.67)	-0.612*** (-18.54)
Observations	4,543	4,543	4,543	4,543
R ²	0.058	0.086	0.120	0.131
Controls	Yes	Yes	Yes	Yes
Contract Province × Year FE	Yes	Yes	Yes	Yes
Officer × Year FE	Yes	Yes	Yes	Yes
Agent × Year FE	Yes	Yes	Yes	Yes

Table 9. Promotion-Driven Rescued Loans and Their Dynamic Loan Outcomes

This table examines how promotion-period rescued loans affect loan outcomes and how these effects evolve over time. The sample is restricted to rescued passenger vehicle loans, defined as applications initially rejected by the machine but subsequently approved by a human reviewer. The dependent variables include: a default indicator equal to one if the loan defaults and zero otherwise; the loan profit ratio, defined as profit divided by initial investment; and Coarse-Group Information Reliance. *Promotion* is an indicator variable equal to one for loans originated during promotion periods and zero otherwise. *Late period* is an indicator equal to one for loans originated in the late post-adoption period and zero for those originated in the early post-adoption period. *Rescued* is an indicator variable equal to one for applications initially rejected by the machine but subsequently approved by a human reviewer, and zero otherwise. All specifications include contract-, borrower-, and vehicle-level controls, as well as contract-province-by-year, officer-by-year, and agent-by-year fixed effects. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle Sample Variables	Passenger					
	Rescued					
	Default	Profit Ratio	CGIR	Default	Profit Ratio	CGIR
	(1)	(2)	(3)	(4)	(5)	(6)
Rescued × Promotion	−0.081*** (−4.09)	0.040*** (3.77)	0.063** (2.49)	−0.162* (−1.79)	0.031 (0.64)	0.257*** (3.63)
Rescued × Promotion × Late Period				0.040 (0.66)	0.009 (0.24)	−0.111** (−2.60)
Promotion × Late Period				0.030 (0.68)	−0.001 (−0.05)	0.030 (1.03)
Rescued × Late Period				−0.065*** (−2.96)	0.009 (0.92)	0.127*** (4.00)
Late Period				0.035*** (4.96)	−0.006 (−1.07)	−0.077*** (−3.52)
Rescued	0.022 (1.17)	−0.013* (−1.89)	−0.055** (−2.04)	0.131** (2.61)	−0.029 (−1.54)	−0.276*** (−4.01)
Promotion	0.014 (0.78)	−0.012 (−0.83)	−0.011 (−0.25)	−0.027 (−0.49)	−0.012 (−0.47)	−0.053 (−0.97)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Contract Province × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Officer × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Agent × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,543	4,543	4,543	4,543	4,543	4,543
<i>R</i> ²	0.098	0.120	0.291	0.100	0.120	0.294

Table 10. Heterogeneity Analysis

This table reports heterogeneity results. The dependent variables include: $\text{Log}(\text{Audit Time})$, $\text{Log}(\text{Text Length})$, Information Density, Fuzzy Dummy, and Textual Similarity. NearCutoff^{30} is an indicator equal to one if the machine score is within 30 points of the cutoff (on the rejected side), and zero otherwise. Late is an indicator equal to one for applications in the late post-adoption period. Heterogeneity measures are constructed as follows: Tenure is the number of months an officer has been in the senior-approval position at the time of FinTech adoption; Past Profit Ratio is the officer's average loan profit ratio among pre-FinTech approvals; Past CGIR is the officer's average CGIR among pre-FinTech approvals; $\text{Vol}(\text{Past Profit Ratio})$ and $\text{Vol}(\text{Past CGIR})$ are the within-officer standard deviations of monthly pre-FinTech profit ratio and CGIR, respectively; COVID Exposure is the city-year share of days classified as medium/high risk. All specifications include contract-, borrower-, and vehicle-level controls, as well as contract-province-by-year, officer-by-year, and agent-by-year fixed effects. Standard errors are two-way clustered at the officer-by-year level, with t -statistics reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle Sample Variables	Passenger				
	Rescued				
	Log(Audit Time)	Log(Text Length)	Info. Density	Fuzzy Dummy	Textual Similarity
	(1)	(2)	(3)	(4)	(5)
Panel A: Tenure (Months on the Position)					
Late \times NearCutoff ³⁰ \times Tenure	-0.061*** (-2.98)	-0.044** (-2.19)	-0.031* (-1.78)	0.023** (2.33)	0.027*** (2.95)
Late \times NearCutoff ³⁰	-0.058** (-2.31)	-0.039* (-1.83)	-0.024 (-1.33)	0.019* (1.76)	0.021** (2.05)
NearCutoff ³⁰ \times Tenure	-0.012 (-0.72)	-0.010 (-0.66)	-0.008 (-0.58)	0.006 (0.78)	0.009 (1.23)
Late \times Tenure	-0.018 (-1.05)	-0.013 (-0.86)	-0.009 (-0.67)	0.007 (0.88)	0.006 (0.78)
Late	-0.072*** (-3.41)	-0.051*** (-2.87)	-0.034** (-2.14)	0.020** (2.12)	0.016 (1.54)
NearCutoff ³⁰	0.090*** (3.06)	0.056** (2.22)	0.039* (1.73)	-0.014 (-1.06)	-0.018* (-1.79)
Panel B: Pre-FinTech Track Record (Past Profit Ratio)					
Late \times NearCutoff ³⁰ \times Past Profit Ratio	-0.032* (-1.84)	-0.028* (-1.75)	-0.015 (-1.07)	0.013* (1.79)	0.018** (2.18)
Late \times NearCutoff ³⁰	-0.049** (-2.01)	-0.035* (-1.72)	-0.021 (-1.22)	0.017* (1.70)	0.023** (2.34)
NearCutoff ³⁰ \times Past Profit Ratio	0.009 (0.62)	0.006 (0.46)	0.003 (0.27)	-0.006 (-0.86)	-0.004 (-0.61)
Late \times Past Profit Ratio	0.014 (0.98)	0.009 (0.71)	0.006 (0.52)	-0.004 (-0.62)	-0.007 (-1.04)
Late	-0.064*** (-3.08)	-0.046*** (-2.74)	-0.029** (-2.00)	0.018** (2.05)	0.020* (1.88)
NearCutoff ³⁰	0.081*** (2.78)	0.049** (2.06)	0.034* (1.70)	-0.017 (-1.32)	-0.016 (-1.56)
Panel C: Pre-FinTech Track Record (Past CGIR)					
Late \times NearCutoff ³⁰ \times Past CGIR	-0.045** (-2.53)	-0.036** (-2.21)	-0.028** (-2.06)	0.020** (2.34)	0.022** (2.20)
Late \times NearCutoff ³⁰	-0.054** (-2.19)	-0.041** (-2.05)	-0.026* (-1.77)	0.021** (2.13)	0.019* (1.90)
NearCutoff ³⁰ \times Past CGIR	-0.016 (-1.05)	-0.019* (-1.67)	-0.021** (-2.02)	0.014* (1.70)	0.010 (1.28)
Late \times Past CGIR	-0.011 (-0.78)	-0.010 (-0.79)	-0.013 (-1.17)	0.010 (1.30)	0.009 (1.15)
Late	-0.070*** (-3.28)	-0.053*** (-3.05)	-0.033** (-2.10)	0.019** (2.10)	0.017 (1.58)
NearCutoff ³⁰	0.088*** (3.01)	0.051** (2.11)	0.041** (2.05)	-0.015 (-1.12)	-0.021** (-2.08)

(Continued)

Vehicle Sample Variables	Passenger				
	Rescued				
	Log(Audit Time)	Log(Text Length)	Info. Density	Fuzzy Dummy	Textual Similarity
	(1)	(2)	(3)	(4)	(5)
Panel D: Ability Uncertainty (Volatility of Past Profit Ratio)					
Late \times NearCutoff ³⁰ \times Vol(Past Profit Ratio)	-0.063*** (-3.27)	-0.052*** (-3.03)	-0.039** (-2.42)	0.026*** (2.82)	0.029*** (3.00)
Late \times NearCutoff ³⁰	-0.051** (-2.12)	-0.037* (-1.82)	-0.024 (-1.41)	0.018* (1.75)	0.021** (2.05)
NearCutoff ³⁰ \times Vol(Past Profit Ratio)	-0.010 (-0.63)	-0.008 (-0.57)	-0.006 (-0.49)	0.005 (0.69)	0.006 (0.88)
Late \times Vol(Past Profit Ratio)	-0.015 (-0.96)	-0.012 (-0.87)	-0.009 (-0.73)	0.006 (0.84)	0.007 (1.02)
Late	-0.062*** (-3.02)	-0.047*** (-2.81)	-0.028** (-1.99)	0.017** (2.04)	0.019* (1.89)
NearCutoff ³⁰	0.083*** (2.86)	0.050** (2.12)	0.036* (1.81)	-0.016 (-1.21)	-0.017* (-1.73)
Panel E: Ability Uncertainty (Volatility of Past CGIR)					
Late \times NearCutoff ³⁰ \times Vol(Past CGIR)	-0.019 (-1.06)	-0.014 (-0.87)	-0.011 (-0.78)	0.009 (1.20)	0.010 (1.25)
Late \times NearCutoff ³⁰	-0.055** (-2.22)	-0.040** (-2.02)	-0.025 (-1.49)	0.020** (2.02)	0.022** (2.15)
NearCutoff ³⁰ \times Vol(Past CGIR)	-0.006 (-0.41)	-0.005 (-0.38)	-0.004 (-0.33)	0.004 (0.58)	0.005 (0.74)
Late \times Vol(Past CGIR)	-0.012 (-0.79)	-0.010 (-0.74)	-0.012 (-1.00)	0.010 (1.49)	0.011 (1.57)
Late	-0.069*** (-3.24)	-0.052*** (-3.02)	-0.033** (-2.12)	0.019** (2.14)	0.018* (1.76)
NearCutoff ³⁰	0.087*** (2.98)	0.053** (2.19)	0.040** (2.03)	-0.015 (-1.12)	-0.020** (-2.00)
Panel F: External Uncertainty Shock (COVID City-Year Exposure)					
Late \times NearCutoff ³⁰ \times COVID Exposure	-0.024 (-1.32)	-0.017 (-1.03)	-0.010 (-0.67)	0.016* (1.84)	0.020** (2.15)
Late \times NearCutoff ³⁰	-0.047** (-1.99)	-0.034* (-1.73)	-0.020 (-1.18)	0.016* (1.69)	0.020** (2.08)
NearCutoff ³⁰ \times COVID Exposure	-0.008 (-0.52)	-0.006 (-0.44)	-0.004 (-0.33)	0.009 (1.28)	0.011* (1.73)
Late \times COVID Exposure	-0.013 (-0.82)	-0.011 (-0.77)	-0.014 (-1.11)	0.012* (1.73)	0.014* (1.75)
Late	-0.060*** (-2.96)	-0.045*** (-2.70)	-0.028** (-1.98)	0.017** (2.02)	0.018* (1.78)
NearCutoff ³⁰	0.079*** (2.71)	0.047** (1.99)	0.032* (1.66)	-0.016 (-1.26)	-0.017* (-1.70)
Controls	Yes	Yes	Yes	Yes	Yes
Contract Province \times Year FE	Yes	Yes	Yes	Yes	Yes
Officer \times Year FE	Yes	Yes	Yes	Yes	Yes
Agent \times Year FE	Yes	Yes	Yes	Yes	Yes
Observations	4,543	4,543	4,543	4,543	4,543

Table 11. Dynamic Effects of Upside Overrides on Loan Outcomes and Borrower Characteristics

This table examines the dynamic effects of human overrides of algorithmic rejections in the credit approval process. The sample consists of all passenger vehicle loan applications with credit scores below 475 that were initially rejected by the machine model. *Upside Override* is an indicator equal to one if an application is rejected by the algorithm but subsequently approved by a human loan officer. *Late Period* is an indicator for the later period (relative to the earlier period), capturing post-adoption dynamics. The dependent variables include a default indicator (equal to one if a loan defaults within 60 days and zero otherwise), the loan profit ratio (defined as profit divided by initial investment), Coarse-Group Information Reliance (CGIR), log(client wage), client workplace status, client house ownership, marital status, presence of children, and local residency. Control variables include contract- and vehicle-level characteristics. For default-related outcomes, we additionally control for borrower characteristics. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle	Passenger					
Sample	Machine rejected (Score < 475)					
Panel A: Loan Outcomes						
Variables	Default	Loan Profit Ratio			CGIR	
	(1)	(2)	(3)	(3)		
Upside Override × Late Period	-0.061*** (-2.79)	0.015 (1.01)		0.197*** (5.33)		
Upside Override	0.060 (1.61)	-0.000 (-0.01)		-0.370*** (-7.65)		
Late Period	0.004 (0.10)	-0.003 (-0.12)		-0.189*** (-4.96)		
Observations	6,721	6,721		6,721		
R^2	0.055	0.138		0.428		
Panel B: Borrower Characteristics						
Variables	Log(Client Wage)	Client Workplace	Client House Ownership	Married Client	Has Children	Local Client
	(1)	(2)	(3)	(4)	(5)	(6)
Upside Override × Late Period	0.210* (1.99)	0.198*** (4.50)	0.252*** (4.60)	-0.005 (-0.11)	0.025 (0.84)	-0.025 (-0.63)
Upside Override	-0.296 (-1.62)	-0.365*** (-5.08)	-0.292*** (-3.28)	0.031 (0.39)	0.027 (0.57)	-0.041 (-0.78)
Late Period	-0.111 (-1.00)	-0.182*** (-3.81)	-0.281*** (-5.37)	-0.012 (-0.24)	-0.028 (-1.01)	0.008 (0.21)
Observations	6,721	6,721	6,721	6,721	6,721	6,721
R^2	0.325	0.191	0.291	0.054	0.084	0.660
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Contract Province × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Officer × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Agent × Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 12. Dynamic Effects of Downside Overrides on Loan Outcomes and Borrower Characteristics

This table examines the dynamic effects of downside overrides in the credit approval process. The sample consists of all passenger vehicle loan applications that were initially approved by the machine model. *Downside Override* is an indicator equal to one if an application is approved by the algorithm but subsequently rejected by a human loan officer. *Late Period* is an indicator for the later period (relative to the earlier period). The dependent variables include a default indicator (equal to one if a loan defaults within 60 days and zero otherwise), the loan profit ratio (defined as profit divided by initial investment), Coarse-Group Information Reliance (CGIR), log(client wage), client workplace status, client house ownership, marital status, presence of children, and local residency. Control variables include contract- and vehicle-level characteristics; for default-related outcomes, borrower-level controls are additionally included. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle	Passenger					
Sample	Machine approved					
Panel A: Loan Outcomes						
Variables	Default	Loan Profit Ratio			CGIR	
	(1)	(2)	(3)	(3)		
Downside Override × Late Period	0.108*** (5.66)	0.050*** (3.94)		0.043 (0.92)		
Downside Override	-0.014 (-0.37)	-0.090*** (-3.91)		-0.023 (-0.31)		
Late Period	-0.020 (-0.91)	0.013 (0.98)		-0.057 (-1.66)		
Observations	10,068	10,068		10,068		
<i>R</i> ²	0.118	0.169		0.247		
Panel B: Borrower Characteristics						
Variables	Log(Client Wage)	Client Workplace	Client House Ownership	Married Client	Has Children	Local Client
	(1)	(2)	(3)	(4)	(5)	(6)
Downside Override × Late Period	-0.260 (-1.67)	0.048 (1.22)	0.139*** (5.01)	0.104*** (2.73)	0.093*** (5.88)	-0.009 (-0.33)
Downside Override	0.200 (0.83)	-0.074 (-1.07)	0.078* (1.92)	-0.352*** (-6.65)	-0.161*** (-5.97)	0.001 (0.02)
Late Period	-0.029 (-0.40)	-0.063*** (-3.94)	-0.029 (-1.15)	-0.002 (-0.07)	-0.048*** (-4.18)	-0.019 (-0.86)
Observations	10,068	10,068	10,068	10,068	10,068	10,068
<i>R</i> ²	0.286	0.106	0.239	0.059	0.074	0.558
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Contract Province × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Officer × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Agent × Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Appendix A: Figures and Tables

Figure A1. Sample of Commercial and Passenger Vehicles

This figure presents representative samples of (a) commercial vehicles and (b) passenger vehicles. Commercial vehicles are designed primarily for the transportation of goods and include buses, trucks, and semi-tractors. Their main users are logistics and transportation drivers. Passenger vehicles are intended for carrying individuals and their personal belongings or small items, and include sedans, sport utility vehicles (SUVs), and similar models. The customer base for passenger vehicles consists predominantly of natural persons.

(a) Commercial Vehicle



(b) Passenger Vehicle



Figure A2. Credit Approval Guidelines before Machine Adoption

This figure illustrates the credit approval guidelines used prior to the adoption of machine-based scoring. Loan approval decisions are based on information contained in the submitted application materials, including borrower characteristics, vehicle information, financing terms, supporting application documents, credit information (such as court enforcement records and credit bureau reports), and the intended use of the vehicle.

汽车融资租赁零售业务信审管理办法

1. 业务/管理流程

1.1 信审管理的主要工作是从收到汽车融资租赁零售业务申请开始,通过信审资料审查方法,经信审业务人员对客户申请的初审及复审,根据客户的综合资质情况出具最终审批意见,审批结果包括通过、拒绝。

1.2 信审人员在初审阶段需要对信审的相关文件及零售业务系统中的相关信息进行完整性、合规性、真实性检查。信审初审人员需通过电话向客户核实身份、车辆及融资条件等信息,并了解客户的工作收入和生活状况,通过电核联系人、核查网络信息等方式对客户提供的信息及申请资料进行交叉验证。

1.3 如初审通过,初审岗提交初审意见并将业务流程提交至信审复审岗;如经初审审查判断需调整融资方案或补充资料的,则由初审岗在零售业务系统中将流程退回至经销商/代理商,待经销商/代理商将信审资料补齐或更正后再次提交信审:如未通过初审审查,初审岗应提交未通过审查意见并将流程提交至复审岗,由复审岗做出拒绝决策,初审岗无直接拒绝权限。复审岗有权对初审意见进行驳回,并由初审岗做重新审查处理。

1.4 信审审批流程结束后,业务流程流转至经销商或代理商处,信审流程结束。

2. 信用审查政策

2.1 基本准则

综合考虑申请当事人基本信息、车辆信息、融资信息、申请材料、信用信息、车辆用途。

2.2 法院执行记录判断准则

2.2.1 关注身份证命中法院失信名单,被执行人的履行性质为“全部未履行”等情况。

2.2.2 关注立案时间距申请日期6个月内,且执行标的 ≥ 1000 元,以“涉诉”拒绝等情况。

2.3 征信报告判断准则

关注商业贷款/个人消费贷款/汽车贷款/信用卡下列情况:

存在非正常终止账户,即贷款账户状态为呆账、核销、强制执行还款、资产处置等;信用卡账户状态为冻结、呆账、核销、止付等;贷款或信用卡出现担保人代偿(D)、以资抵债(Z)、结束(G)等;

近2年内逾期出现过“3”及以上的或最近2年内出现过两次“2”及以上的;

贷款、担保贷款出现次级、可疑、损失。

Figure A3. Comparison of Decision-Making Principles between Human Credit Officers and the Machine-Based Scoring System

This figure compares the decision-making principles of human credit officers and the machine-based scoring system. Human credit officers rely primarily on experience-based judgment and discretionary evaluation of application materials, while the machine-based system applies data-driven algorithms to assess borrower risk using a standardized and high-dimensional information set.

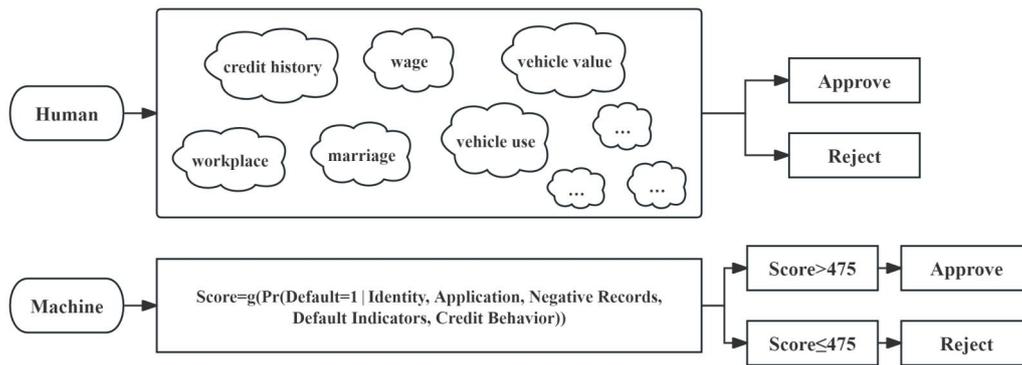


Figure A4. Parallel Trends of Default Rate and CGIR

This figure presents the parallel trends in default rate and Coarse-Group Information Reliance (CGIR) using the full sample of commercial and passenger vehicle loan applications for the staggered DID analysis.

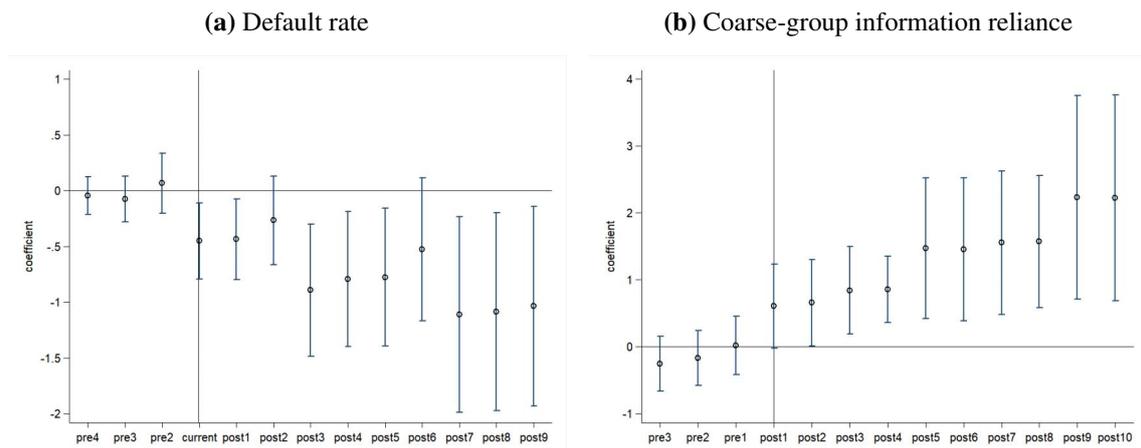


Table A1. Default Handling Measures by Type of Default

This table shows that default handling escalates based on severity and frequency. Regular methods are applied initially, while deep and strong interventions are reserved for repeat or serious delinquencies.

Disposal Measures	First-time Default	Early-stage Default (< 60 days)	Non-first-time / Non-early Default
<i>Regular Measures:</i>			
Smartphone reminder	✓	✓	
Manual phone reminder	✓	✓	✓
<i>Deep Measures:</i>			
Report to credit registry	✓	✓	✓
Lock vehicle via GPS	✓	✓	✓
Send letter		✓	✓
<i>Strong Measures:</i>			
Door-to-door reminder		✓	✓
Vehicle repossession		✓	✓
Legal proceedings			✓

Table A2. Information Set Used by the FinTech Scoring System

This table describes the information set utilized by the FinTech scoring system. The machine-generated score is based on a high-dimensional set of customer data, including identity verification, application information, negative records, default-related indicators, and credit behavior information.

Type	Information Content
Identity information	Verification of the applicant’s identity through biometric data (e.g., facial recognition) and bank card authentication.
Application information	Borrower characteristics, vehicle information, financing terms, supporting application documents, and intended use of the vehicle.
Negative information	Records of adverse financial conduct in previous financial institution interactions. Judicial records, including case filing status, enforcement actions, and listings of dishonest behavior.
Default information	Detection of multiple loan applications submitted by the applicant across institutions. Assessment of the applicant’s repayment burden based on financial and behavioral indicators. Identification of potential risks associated with gang-related fraud or intermediary exploitation.
Credit information	Evaluation of the applicant’s consumption habits, financial behavior, and repayment stability based on comprehensive scoring models.

Table A3. Top Predictors Before and After FinTech Adoption

This table reports the top features ranked by machine-learning-based model importance in the pre-adoption and post-adoption periods. Importance values are computed from machine-learning-based predictions for each feature within period and are not necessarily comparable across periods.

Vehicle		Passenger + Commercial		
Sample	Pre-FinTech		Post-FinTech	
Rank	Feature	Importance	Feature	Importance
1	Client House Ownership	0.174	Client House Ownership	0.076
2	Client Workplace	0.169	Client Workplace	0.044
3	New Vehicle	0.014	Log(Contract Duration)	0.037
4	Local Client	0.008	Married Client	0.032
5	Married Client	0.006	Local Client	0.025
6	Contract Interest Rate	0.003	Contract Interest Rate	0.023
7	Log(Contract Duration)	0.003	Log(Client Wage)	0.015
8	Business Use	0.002	Log(Contract Financing)	0.008
9	Log(Contract Financing)	0.002	Has Children	0.008
10	Down Payment Ratio	0.002	Down Payment Ratio	0.007

Table A4. Summary Statistics by Year and by Vehicle Type

This table reports the number of approved applications distinguishing between the pre- and post-FinTech adoption periods. The year 2020 is split into two subperiods: H1 and H2, corresponding to the staggered rollout of the FinTech credit scoring system. For commercial vehicles, the post-adoption period begins in March 2020; for passenger vehicles, it begins in July 2020.

Period	Year	Commercial vehicle		Passenger vehicle		Total	
		Number	Percentage	Number	Percentage	Number	Percentage
Pre-FinTech	2018	115	0.08%	143	0.70%	258	0.16%
	2019	6,276	4.43%	2,613	12.87%	8,889	5.49%
	2020 (H1)	22,789	16.10%	2,577	12.69%	25,366	15.68%
	Sum	29,180	20.62%	5,333	26.26%	34,513	21.33%
Post-FinTech	2020 (H2)	40,430	28.57%	4,636	22.83%	45,066	27.85%
	2021	57,437	40.59%	9,623	47.39%	67,060	41.44%
	2022	4,195	2.96%	338	1.66%	4,533	2.80%
	2023	10,270	7.26%	376	1.85%	10,646	6.58%
	Sum	112,332	79.38%	14,973	73.74%	127,305	78.67%
Total	–	141,512	100.00%	20,306	100.00%	161,818	100.00%

Table A5. Effects of FinTech Credit Scoring Adoption on Early Default Probability

This table examines the effects of adopting the FinTech credit scoring system on early default probability. The dependent variables are indicator variables equal to one if a loan defaults within 6, 12, 18, 24, 30, or 36 months after loan origination, respectively. The treatment variable, *Treat*, equals one for loan applications submitted after the implementation of the FinTech scoring system—March 2020 for commercial vehicles and July 2020 for passenger vehicles—and zero otherwise. Control variables include contract-level, borrower-level, and vehicle-level characteristics. The samples include loans with at least 36 months of data. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. Coefficients marked with *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle Variable	Passenger + Commercial					
	DEF ⁶	DEF ¹²	DEF ¹⁸	DEF ²⁴	DEF ³⁰	DEF ³⁶
	(1)	(2)	(3)	(4)	(5)	(6)
Treat	-0.004 (-0.86)	0.003 (0.25)	0.006 (0.42)	0.021 (0.89)	0.025 (1.20)	-0.006 (-0.28)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Contract Province × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Officer × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Agent × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	37,267	37,267	37,267	37,267	37,267	37,267
R ²	0.183	0.206	0.237	0.246	0.247	0.213

Table A6. Effects of FinTech Adoption and CGIR on Loan Contract Terms

This table reports regression results examining the effects of FinTech adoption and Coarse-Group Information Reliance (CGIR) on loan-to-value (LTV) and contract interest rates. Control variables include contract-level and vehicle-level characteristics. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. Coefficients marked with *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle Variables	Passenger + Commercial	
	LTV	Contract Interest Rate
Treat × CGIR	0.009*** (9.82)	0.180 (1.32)
Treat	0.004*** (4.64)	0.173 (1.14)
CGIR	0.003*** (4.12)	-0.233* (-1.77)
Controls	Yes	Yes
Contract Province × Year FE	Yes	Yes
Officer × Year FE	Yes	Yes
Agent × Year FE	Yes	Yes
Observations	134,467	134,467
R ²	0.945	0.739

Table A7. Placebo Test: Fictitious FinTech Scoring Adoption Dates

This table reports placebo (falsification) tests using fictitious adoption dates of the FinTech credit scoring system. The dependent variables include a default indicator (equals one if a loan defaults and zero otherwise), the loan profit ratio (profit divided by initial investment), and the Coarse-Group Information Reliance index (CGIR). The placebo treatment variable, *Treat*, equals one for loan applications submitted after the *fictitious* implementation dates (reported in the column headers) and zero otherwise. All specifications include contract-, borrower-, and vehicle-level controls, as well as contract-province-by-year, officer-by-year, and agent-by-year fixed effects. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle	Passenger + Commercial					
	Commercial Aug. 2021, Passenger Jul. 2022			Commercial May 2022, Passenger Feb. 2022		
Fictitious Adoption Date	Default	Loan Profit Ratio	CGIR	Default	Loan Profit Ratio	CGIR
Variable	(1)	(2)	(3)	(4)	(5)	(6)
Treat	−0.006 (−0.74)	0.003 (0.51)	−0.005 (−0.63)	0.004 (0.58)	−0.002 (−0.36)	0.003 (0.49)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Contract Province × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Officer × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Agent × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	134,467	99,013	105,492	134,467	99,013	105,492
R^2	0.244	0.305	0.467	0.244	0.305	0.467

Table A8. Effects of FinTech Credit Scoring Adoption on Default Probability and Loan Profit: Passenger Vehicle Sample

This table examines changes in loan outcomes before and after the adoption of the FinTech scoring system, using the sample of passenger vehicle contracts. The dependent variables include a default indicator, which equals one if a loan defaults and zero otherwise, and the loan profit ratio, defined as profit divided by initial investment. Treatment variable equals 1 if the contract is issued in a province with lower levels of financial development, and 0 otherwise. We proxy financial development using three measures: digital finance (*Treat Digi*), total social financing (*Treat Social Fin*), and stock market activity (*Treat Stock*). Control variables include contract-level, borrower-level, and vehicle-level characteristics. “Uncensored” loans refer to loans whose original maturities do not extend beyond the end of the sample period, while “completed” loans are those that are fully completed within the sample period. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. Coefficients marked with *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle Variable Sample	Passenger					
	Default			Loan Profit Ratio		
	All	Uncensored	Completed	All	Uncensored	Completed
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: <i>Treat Digi</i>						
Treat Digi × Post	-0.441*** (-21.84)	-0.449*** (-22.61)	-0.426*** (-22.01)	0.059*** (4.57)	0.068*** (5.89)	0.057*** (5.05)
Observations	19,209	12,355	13,515	19,209	12,355	13,515
<i>R</i> ²	0.188	0.223	0.291	0.121	0.112	0.205
Panel B: <i>Treat Social Fin</i>						
Treat Social Fin × Post	-0.441*** (-21.84)	-0.449*** (-22.61)	-0.426*** (-22.01)	0.059*** (4.57)	0.068*** (5.89)	0.057*** (5.05)
Observations	19,209	12,355	13,515	19,209	12,355	13,515
<i>R</i> ²	0.184	0.218	0.287	0.121	0.111	0.205
Panel C: <i>Treat Stock</i>						
Treat Stock × Post	-0.441*** (-21.84)	-0.449*** (-22.61)	-0.426*** (-22.01)	0.059*** (4.57)	0.068*** (5.89)	0.057*** (5.05)
Observations	19,209	12,355	13,515	19,209	12,355	13,515
<i>R</i> ²	0.184	0.217	0.286	0.121	0.110	0.204
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Contract Province × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Officer × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Agent × Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A9. Effects of FinTech Credit Scoring Adoption on Coarse-Group Information Reliance: Passenger Vehicle Sample

This table examines the effects of adopting the FinTech credit scoring system on coarse-group information reliance in passenger vehicle sample. The dependent variables include Coarse-Group Information Reliance (CGIR), log(client wage), client workplace, client house ownership, married client, has children, and local client. Treatment variable equals 1 if the contract is issued in a province with lower levels of financial development, and 0 otherwise. We proxy financial development using three measures: digital finance (*Treat Digi*), total social financing (*Treat Social Fin*), and stock market activity (*Treat Stock*). Control variables include contract-level and vehicle-level characteristics. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. Coefficients marked with *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle	Passenger						
Variable	CGIR	Log(Client Wage)	Client Workplace	Client House Ownership	Married Client	Has Children	Local Client
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: <i>Treat Digi</i>							
Treat Digi × Post	0.364*** (19.12)	0.283*** (6.67)	0.242*** (10.06)	0.152*** (4.95)	0.008 (0.37)	0.209*** (10.85)	0.060*** (3.10)
Observations	19,209	19,209	19,209	19,209	19,209	19,209	19,209
<i>R</i> ²	0.288	0.330	0.140	0.229	0.035	0.132	0.642
Panel B: <i>Treat Social Fin</i>							
Treat Social Fin × Post	0.342*** (11.61)	0.217*** (3.11)	0.243*** (16.79)	0.124*** (4.48)	0.031 (1.48)	0.179*** (8.03)	0.056*** (3.28)
Observations	19,209	19,209	19,209	19,209	19,209	19,209	19,209
<i>R</i> ²	0.288	0.330	0.141	0.228	0.035	0.131	0.642
Panel C: <i>Treat Stock</i>							
Treat Stock × Post	0.336*** (11.14)	0.170** (2.22)	0.253*** (19.22)	0.120*** (3.91)	0.040** (2.02)	0.175*** (6.73)	0.047*** (3.48)
Observations	19,209	19,209	19,209	19,209	19,209	19,209	19,209
<i>R</i> ²	0.287	0.329	0.141	0.228	0.035	0.130	0.642
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Contract Province × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Officer × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Agent × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table A10. Effects of FinTech Adoption and CGIR on Loan Contract Terms: Passenger Vehicle Sample

This table reports regression results examining the effects of FinTech adoption and Coarse-Group Information Reliance on loan-to-value (LTV) and contract interest rates in passenger vehicle sample. Treatment variable equals 1 if the contract is issued in a province with lower levels of financial development, and 0 otherwise. We proxy financial development using three measures: digital finance (*Treat Digi*), total social financing (*Treat Social Fin*), and stock market activity (*Treat Stock*). Control variables include contract-level and vehicle-level characteristics. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. Coefficients marked with *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle Variable	Passenger					
	LTV			Contract Interest Rate		
	Treat Digi	Treat Social Fin	Treat Stock	Treat Digi	Treat Social Fin	Treat Stock
	(1)	(2)	(3)	(4)	(5)	(6)
Treat × Post × CGIR	0.012*** (5.44)	0.007*** (2.95)	0.005* (1.96)	0.218 (1.16)	0.311 (1.23)	0.347 (1.49)
Treat × Post	0.008*** (6.69)	0.006*** (4.38)	0.005*** (2.66)	-1.248*** (-10.90)	-0.986*** (-9.31)	-0.975*** (-9.95)
Treat × CGIR	-0.002 (-1.44)	-0.001 (-0.56)	-0.000 (-0.14)	-0.391*** (-2.75)	-0.462*** (-2.95)	-0.503*** (-3.10)
Post × CGIR	-0.000 (-0.30)	0.001 (0.67)	0.002 (1.14)	-0.266 (-1.36)	-0.234 (-0.93)	-0.202 (-0.87)
CGIR	0.003*** (4.91)	0.004*** (4.72)	0.003*** (4.40)	0.153* (1.85)	0.134 (1.59)	0.108 (1.28)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Contract Province × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Officer × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Agent × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	19,209	19,209	19,209	19,209	19,209	19,209
<i>R</i> ²	0.968	0.967	0.967	0.840	0.840	0.840

Table A11. Dynamics of Rescued Loan Concentration around the FinTech Score Cutoff

This table examines how the concentration of approved passenger vehicle loans around the FinTech score cutoff evolves over time following the adoption of the FinTech credit scoring system. The sample is restricted to approved passenger vehicle loans in the post-adoption period. The dependent variables, *Near cutoff*¹⁰, *Near cutoff*²⁰, and *Near cutoff*³⁰, are indicator variables equal to one if a loan's FinTech score falls within increasingly narrow intervals below the cutoff of 475, specifically [465, 475], [455, 475], and [445, 475], respectively. We additionally consider the logarithm of the distance between the FinTech score and the cutoff, *Log(Score distance)*, as an alternative continuous measure of proximity to the cutoff. The key explanatory variable, *Late period*, is an indicator equal to one for loans originated in the late post-adoption period and zero for those originated in the early post-adoption period. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. Coefficients marked with *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle Sample Variables	Passenger			
	Rescued (Fintech Score < 475)			
	Near cutoff ¹⁰	Near cutoff ²⁰	Near cutoff ³⁰	Log(score distance)
	(1)	(2)	(3)	(4)
Late period	0.057*** (7.47)	0.121*** (11.95)	0.225*** (18.89)	-0.621*** (-20.54)
Controls	Yes	Yes	Yes	Yes
Contract Province × Year FE	Yes	Yes	Yes	Yes
Officer × Year FE	Yes	Yes	Yes	Yes
Agent × Year FE	Yes	Yes	Yes	Yes
Observations	4,543	4,543	4,543	4,543
<i>R</i> ²	0.056	0.084	0.119	0.134

Table A12. Keyword Dictionaries for Information Density

This table reports the keyword dictionaries used to construct the two subcomponents of the Information Density measure: Factual Entity Density and Feature-Referencing Density. Keywords are translated into English for exposition, while the original texts are in Chinese. Factual entity density captures references to concrete, case-specific information such as time, location, identity, and verification outcomes. Feature-referencing density captures explicit references to borrower characteristics, credit history, loan attributes, and risk indicators.

Category	Representative Keywords (English Translation)
<i>Factual Entity Density</i>	local resident, non-local resident, residing locally, recently relocated, many years, renting, rental housing, self-built housing, property ownership, homeowner, no property, under parents' ownership, registered property, residence, place of residence, household, family members, household size, married, unmarried, no children, with children, child, birth, age, national ID, employment, workplace, employer, employer verification, verification completed, phone verification, non-cooperative, cooperative, contacted, contacted spouse, spouse, self-reported, applicant self-report, application form, form completion, non-existent individual, false information, unverifiable information, online verification, registration records, self-employed, clothing shop, construction, utilities, subcontracting, project-based work, business license, taxi driver, ride-hailing driver, hotel employment, agricultural bureau, emergency contact, contact person, relative, proxy purchase, vehicle purchased by proxy, purchase purpose, abnormal purchase purpose, vehicle user, unclear relationship
<i>Feature-Referencing Density</i>	credit report, credit report indicates, credit history, no credit record, historical delinquency, delinquency, consecutive delinquency, maximum delinquency, settled delinquency, poor credit, credit card, prior approval rejected, loan, agricultural loan, auto loan, loan status, watchlist, outstanding balance, remaining four installments, number of installments, monthly payment, bank statements, transaction records, account inflow, average monthly inflow, monthly average, repayment burden, repayment pressure, liabilities, high leverage, income, wage, third-party risk score, enforcement record, court enforcement, enforcement court, filing date, case number, risk type, risk hit, blacklist, dishonest debtor, fraud

Table A13. Keyword Dictionaries for Fuzzy Expressions in Approval Opinions

This table reports the dictionary of fuzzy expressions used to construct the Fuzziness Share measure. Keywords are translated into English for exposition, while the original texts are in Chinese. Fuzzy expressions refer to vague, non-specific, and easily defensible phrases that lack explicit numerical, factual, or feature-based justification and can be flexibly reused across cases.

Category	Fuzzy Expressions (English Translation)
<i>Vague evaluative judgments</i>	close to the standard, close to system requirements, basically compliant, compliant with basic requirements, overall acceptable, generally acceptable, acceptable
<i>Non-committal risk assessments</i>	risk controllable, risk acceptable, risk not high, relatively low risk, overall risk controllable
<i>Defensive summary statements</i>	comprehensive evaluation, comprehensive judgment, comprehensive consideration, overall judgment, evaluation suggests
<i>Low-specificity qualifiers</i>	acceptable, slightly low, marginally low, relatively stable, basically stable

Table A14. Robustness Check: Dynamic Changes in Approval Text Characteristics near the FinTech Score Cutoff

This table examines the dynamic evolution of approval text characteristics following the adoption of the FinTech credit scoring system for passenger vehicle loans. The sample is restricted to approved applications only. Textual outcomes include the logarithm of text length, information density, a dummy indicator for the presence of fuzzy or vague expressions, and the average semantic similarity across approval opinions. *Late period*, is an indicator equal to one for loans originated in the late post-adoption period and zero for those originated in the early post-adoption period. *Near cutoff*¹⁰ and *Near cutoff*²⁰ indicate whether the FinTech score falls within [465,475] and [455,475], respectively. All specifications include contract-, borrower-, and vehicle-level controls, as well as contract-province-by-year, officer-by-year, and agent-by-year fixed effects. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle	Passenger			
Sample	Approved			
Variables	Log(Text Length)	Info. Density	Fuzzy Dummy	Textual Similarity
Panel A: Near Cutoff¹⁰				
Late Period × Near cutoff ¹⁰	−0.173*** (−9.31)	−0.003** (−2.26)	0.016*** (15.82)	0.016 (0.75)
Late Period	−0.052*** (−8.44)	−0.001 (−1.58)	0.004*** (9.33)	0.029** (2.18)
Near cutoff ¹⁰	0.096*** (3.03)	0.002 (1.05)	−0.012*** (−9.75)	0.007 (0.21)
Observations	4,543	4,543	4,543	4,523
R ²	0.055	0.023	0.128	0.181
Panel B: Near Cutoff²⁰				
Late Period × Near Cutoff ²⁰	−0.134*** (−7.44)	−0.002* (−1.81)	0.013*** (13.27)	0.010 (0.44)
Late Period	−0.036*** (−6.17)	−0.000 (−1.02)	0.003*** (7.77)	0.027* (1.88)
Near Cutoff ²⁰	0.124*** (4.42)	0.002 (1.14)	−0.013*** (−14.49)	0.009 (0.26)
Observations	4,543	4,543	4,543	4,523
R ²	0.072	0.024	0.175	0.185
Controls	Yes	Yes	Yes	Yes
Contract Province × Year FE	Yes	Yes	Yes	Yes
Officer × Year FE	Yes	Yes	Yes	Yes
Agent × Year FE	Yes	Yes	Yes	Yes

Table A15. Effects of Upside Overrides on Loan Outcomes and Borrower Characteristics

This table examines the effects of human overrides of algorithmic rejections in the credit approval process. The sample consists of all passenger vehicle loan applications with credit scores below 475 that were initially rejected by the machine model. *Upside Override* is an indicator equal to one if an application is rejected by the algorithm but subsequently approved by a human loan officer. The dependent variables include a default indicator (equal to one if a loan defaults within 60 days and zero otherwise), loan profit ratio (defined as profit divided by initial investment), Coarse-Group Information Reliance (CGIR), log(client wage), client workplace status, client house ownership, marital status, presence of children, and local residency. Control variables include contract- and vehicle-level characteristics. For default-related outcomes, we additionally control for borrower characteristics. Standard errors are two-way clustered at the officer-by-year level, with *t*-statistics reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle	Passenger					
Sample	Machine rejected (Score < 475)					
<i>Panel A: Loan Outcomes</i>						
Variables	Default	Loan Profit Ratio			CGIR	
	(1)	(2)	(2)	(2)	(3)	
Upside Override	-0.032** (-2.14)	0.022** (2.51)			-0.089*** (-3.09)	
Observations	6,721	6,721			6,721	
R^2	0.054	0.138			0.424	
<i>Panel B: Borrower Characteristics</i>						
Variables	Log(Client Wage)	Client Workplace	Client House Ownership	Married Client	Has Children	Local Client
	(1)	(2)	(3)	(4)	(5)	(6)
Upside Override	0.013 (0.20)	-0.081*** (-3.06)	0.064 (1.32)	0.023 (0.85)	0.063*** (4.01)	-0.078*** (-5.65)
Observations	6,721	6,721	6,721	6,721	6,721	6,721
R^2	0.323	0.186	0.286	0.054	0.083	0.660
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Contract Province × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Officer × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Agent × Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A16. Effects of Downside Overrides on Loan Outcomes and Borrower Characteristics

This table examines the effects of downside overrides in the credit approval process. The sample consists of all passenger vehicle loan applications that were initially approved by the machine model. *Downside Override* is an indicator equal to one if an application is approved by the algorithm but subsequently rejected by a human loan officer. The dependent variables include a default indicator (equal to one if a loan defaults within 60 days and zero otherwise), the loan profit ratio (defined as profit divided by initial investment), Coarse-Group Information Reliance (CGIR), $\log(\text{client wage})$, client workplace status, client house ownership, marital status, presence of children, and local residency. Control variables include contract- and vehicle-level characteristics. For default-related outcomes, borrower-level controls are additionally included. Standard errors are two-way clustered at the officer-by-year level, with t -statistics reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Vehicle	Passenger					
Sample	Machine approved (Score>475)					
<i>Panel A: Loan Outcomes</i>						
Variables	Default	Loan Profit Ratio			CGIR	
	(1)	(2)	(3)	(4)	(5)	(6)
Downside Override	0.138*** (6.59)	-0.023*** (-2.77)			0.043* (1.77)	
Observations	10,068	10,068			10,068	
R^2	0.117	0.168			0.246	
<i>Panel B: Borrower Characteristics</i>						
Variables	Log(Client Wage)	Client Workplace	Client House Ownership	Married Client	Has Children	Local Client
	(1)	(2)	(3)	(4)	(5)	(6)
Downside Override	-0.155** (-2.19)	-0.000 (-0.00)	0.274*** (11.83)	-0.208*** (-9.44)	-0.027* (-1.77)	-0.009 (-0.62)
Observations	10,068	10,068	10,068	10,068	10,068	10,068
R^2	0.285	0.105	0.237	0.059	0.072	0.558
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Contract Province \times Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Officer \times Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Agent \times Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Appendix B: Credit Effects of FinTech Score Adoption: Single DiD Evidence

Although the FinTech credit scoring system is introduced simultaneously for all passenger vehicle applicants, its impact is likely heterogeneous across regions. In particular, customers from less financially developed provinces are expected to be more affected by the algorithmic shift, as financial intermediation in these regions tends to be relationship-based and constrained by limited access to formal credit infrastructure. In such settings, creditworthiness assessments often rely on soft information—such as personal reputation, local ties, or subjective judgment by loan officers—that is difficult to digitize. The introduction of standardized, data-driven scoring systems may thus disproportionately affect these borrowers by replacing informal, context-specific signals with uniform statistical rules that may not fully capture their credit potential.

This heterogeneity is evaluated by constructing proxies for the financial development level of each contract province using several measures. Our primary measure is the level of digital financial inclusion (*Treat Digi*), sourced from the Digital Finance Index published by the Institute of Digital Finance at Peking University.¹⁰ For each loan application, the indicator *Treat Digi* takes the value 1 if the contract province’s digital finance index, measured in the year prior to FinTech adoption, falls below the sample mean, and 0 otherwise. Provinces below the mean are considered the treatment group (less digitally developed), and those above the mean serve as the control group. For robustness, we consider two additional proxies for financial development, both obtained from the WIND database.¹¹ The first is the ratio of total social financing to provincial GDP (*Treat Social Fin*), which captures the aggregate volume of funds received by the real economy—such as loans, bonds, and equity financing—from the financial system. The second is the ratio of stock market capitalization to provincial GDP (*Treat Stock*), which reflects local capital market activity. For each indicator, we divide provinces into treatment and control groups based on whether the value falls below or above the sample mean measured in the year prior to FinTech adoption, respectively.

To examine how the introduction of machine-based credit scoring affects borrowers differentially depending on the financial development of their local credit environment, we estimate the following regression model:

$$Y_{k,t} = \beta_0 + \beta_1 (Treatment_p \times Post_t) + \mathbf{X}_{k,t} + \phi_p + \delta_{j,t} + \mu_{a,t} + \varepsilon_{k,t}. \quad (3)$$

where $Treatment_p$ is an indicator variable equal to 1 if the loan contract is from province p , classified as financially underdeveloped based on the measures described above, and 0 otherwise. The variable $Post_t$ takes the value 1 for loan applications submitted on or after July 2020, and 0 otherwise. The control vector $\mathbf{X}_{k,t}$ is defined as in Equation (1), including contract-, vehicle-, and applicant-level

¹⁰<https://idf.pku.edu.cn/>.

¹¹<https://www.wind.com.cn/>.

characteristics. The specification includes province fixed effects (ϕ_p), credit-officer-by-year fixed effects ($\delta_{j,t}$), and dealership-by-year fixed effects ($\mu_{a,t}$), which jointly control for time-invariant regional characteristics, officer-specific heterogeneity that varies over time, and agent-level shocks. Standard errors are clustered at the credit-officer–year level.

Table A8 reports the single difference-in-differences estimates for the passenger vehicle sample. Panel A, Panels B, and C correspond to the three alternative measures of regional financial development. Columns (1)–(3) examine default outcomes. Across all three measures, we find that the interaction term is negative and highly significant, indicating that the introduction of FinTech credit scoring leads to a substantially larger reduction in default rates in financially underdeveloped provinces. The effect is robust when restricting the sample to uncensored loans and to completed loans, suggesting that the improvement in credit performance is not driven by differential censoring or contract maturity. Columns (4)–(6) turn to loan profitability. We find that FinTech adoption is associated with significantly higher loan profit ratios in less financially developed provinces, with economically similar magnitudes across the full, uncensored, and completed samples. Taken together, these results indicate that the algorithmic scoring system generates larger efficiency gains precisely in regions where pre-adoption credit allocation relied more heavily on relationship-based or soft-information-intensive screening.

Table A9 examines the fairness implications of FinTech adoption across regions with different levels of financial development. Column (1) shows that the interaction is positive and highly significant for the composite CGIR measure, indicating that the increase in reliance on coarse-group information is more pronounced in financially underdeveloped provinces. Columns (2) through (7) further decompose this effect and show that FinTech adoption significantly increases the weight placed on $\log(\text{Client Wage})$, *Client Workplace*, *Client House Ownership*, and *Has Children*, while the effect on *Married Client* is weaker and less precisely estimated. In contrast, reliance on *Local Client* declines after FinTech adoption, consistent with the replacement of locally embedded, relationship-based screening with standardized statistical rules in less developed credit markets.

Finally, Table A10 explores whether heightened reliance on coarse-group information translates into changes in contract terms. We find that the triple interaction is positive and statistically significant for loan-to-value (LTV) ratios across all three measures of financial development, indicating that borrowers whose approval decisions rely more heavily on coarse-group information receive more permissive leverage after FinTech adoption. By contrast, we find no statistically significant effects on contract interest rates. These findings suggest that in financially underdeveloped regions, the algorithmic shift not only reshapes approval decisions, but also reallocates credit conditions along the intensive margin.

Overall, the single DiD results reinforce the main findings from the staggered DiD analysis: FinTech credit scoring improves efficiency by reducing default and increasing profitability, but does so by amplifying reliance on coarse-group information, particularly in environments where traditional credit assessment relied more heavily on soft and relationship-based information.