The Value of Man in AI + Man: Field Evidence from Small Business Lending*

Xia Chen
Singapore Management University

Qiang Cheng
Singapore Management University

Guanqun Dou
Fudan University

Shuai Shao
Zhejiang University

April 2026

Abstract

We investigate the value of human loan officers in an AI + Man loan decision model, where loan officers decide on loan application outcomes based on recommendations from an AI model. Using proprietary loan-level data from a lending company that provides loans to small businesses, we find that the incremental contribution of loan officers to loan profits beyond AI recommendations increases with the soft information collected and used by loan officers and decreases with agency issues, algorithm aversion, and mental fatigue of loan officers. Further analyses indicate that using both the quantity and quality of approved loans to evaluate loan officers enhances the value of soft information, mitigates the adverse impact of agency issues, but exacerbates the adverse effect of algorithm aversion and mental fatigue. Our paper contributes to the literature by shedding light on the conditions under which humans can create or destroy value in an AI + Man decision model.

Keywords: artificial intelligence (AI); lending; soft information; agency issues; algorithm aversion; mental fatigue; FinTech

JEL codes: G2, G21, G32, O33

* We wish to thank Asad Kausar, Miao Liu, Alvis Lo, Tharindra Ranasinghe, and workshop participants at American University for their helpful comments. Chen and Cheng acknowledge funding from the Lee Kong Chian Professorship at Singapore Management University. The authors of this paper have no conflicts of interest related to this research. Please contact Xia Chen at xcheng@smu.edu.sg, Qiang Cheng at qcheng@smu.edu.sg, Guanqun Dou at gqundou@163.com, and Shuai Shao at shuaishao@zju.edu.cn for comments.

1 Introduction

In this paper, we examine the role played by human loan officers in an "AI + Man" decision model for small business lending, using proprietary data from a FinTech company. The company uses an artificial intelligence (AI) model to process all loan applications, and loan officers then make final loan approval decisions based on AI's recommendations. The deviation of the final decisions from AI's recommendations captures the value contributed by loan officers. Our aim is to understand the conditions under which loan officers can provide incremental value over AI's recommendations.

The motivation for the research question is three-fold. First, recent advancements in AI have led to substantial investments by both corporations and governments (Agrawal, Gans, and Goldfarb 2019; Babina et al. 2024). For example, Gartner predicts that investments in AI will reach $1.5 trillion in 2025.[1] Many organizations have incorporated AI-driven decision processes and, in many cases, have replaced human decision-making with automated systems (e.g., Acemoglu et al. 2022; Coleman et al. 2022). These developments raise important questions regarding the continued relevance and value of human judgment in environments increasingly mediated by AI. Some commentators have even suggested that human input is going to be redundant for knowledge-based tasks.[2] Examining the value of humans in AI-assisted decisions is thus a timely and economically significant topic.

Second, our research question is particularly timely because an increasing number of companies are providing AI assistance to their workers. For example, a McKinsey report finds that that more than 76% of corporate executives expect workers to use AI tools for more

[1] See details at https://www.gartner.com/en/newsroom/press-releases/2025-09-17-gartner-says-worldwide-ai-spending-will-total-1-point-5-trillion-in-2025.
[2] See "The Most Joyless Tech Revolution Ever: AI Is Making Us Rich and Unhappy" (https://www.wsj.com/tech/ai/the-most-joyless-tech-revolution-ever-ai-is-making-us-rich-and-unhappy-6b7116a3).

than 30% of their daily tasks in the next five years.[3] Prior research provides suggestive evidence that human expertise remains valuable when decisions involve soft information that is difficult to codify (e.g., Brynjolfsson, Mitchell, and Rock 2018; Cao et al. 2024; Jansen et al. 2025). A systematic examination of the conditions under which humans contribute positively or negatively when assisted by AI can improve the design of AI–human collaboration frameworks.

Third, the rapid expansion of FinTech platforms and the adoption of algorithmic credit scoring have transformed lending practices (Fuster et al. 2019; Sutherland 2020; Gopal and Schnabl 2022). For example, Fuster et al. (2019) find that compared with traditional lenders, FinTech lenders process mortgage applications faster without sacrificing risk control. Yet many financial institutions continue to rely on loan officers. Thus, understanding the role of humans in an AI-assisted lending model can inform FinTech companies and banks in staffing strategies in an era of accelerating technological integration (Kleinberg et al. 2018; McKinsey Analytics 2025).

To investigate the value of humans in an AI + Man decision model, we exploit the proprietary data from a lending company in China (referred to as the focal company). The company provides loans to small business owners. Starting from the second half of 2023, the company adopts an AI + Man hybrid decision model. Upon receiving a loan application, the company's proprietary AI model collects and analyzes extensive information about the borrower. The AI model then provides a recommendation on whether to approve the loan application, loan amount, interest rate, and maturity term.[4] Based on the AI recommendation, a loan officer conducts a short online interview with the borrower and the customer officer who handles the loan application and then makes the final decision on whether to approve the

[3] See details at https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/superagency-in-the-workplace-empowering-people-to-unlock-ais-full-potential-at-work#/.
[4] The AI model provides recommendations on loan amount, interest rate, and maturity term even when it recommends denial of loan applications.

loan, and if so, the credit amount, interest rate, and term.

Using the loans approved in 2024 by the focal company, we calculate the loan profit based on loan officers' decisions and the hypothetical profit based on AI's recommendations. The difference between the two, scaled by the amount dispensed to the borrowers, referred to as the deviation in profit, captures the value contributed by human loan officers over AI in the AI + Man hybrid decision model.

We start by investigating whether on average human loan officers contribute value beyond AI recommendations. We find that the average deviation in profit is positive, suggesting that on average human loan officers contribute positively to loan profits over AI. In addition, the variation of the deviation in profit is large, with a standard deviation of about 6%. The deviation in profit is primarily driven by the difference in loan amount and to a lesser extent by the differences in interest rate and maturity term. The large variation of the deviation in profit suggests that it is important to understand the *conditions* under which loan officers contribute positively or negatively.

Based on prior research, we identify three dimensions where loan officers' contributions might vary. First, we focus on the soft information collected and used by loan officers. Prior literature highlights the importance of soft information in lending decisions (e.g., Campbell, Loumioti, and Whittenberg-Moerman 2019). Taking advantage of the comments made by loan officers when making decisions, we code the length of the comments and collect the comments regarding borrowers' business prospects, whether borrowers will receive support from their family members, explanations of lawsuits or cash flow anomalies of borrowers, and whether borrowers are recommended by existing clients or acquaintances. Using these proxies, we find that the use of soft information is positively correlated with the value contributed by loan officers. The effect ranges from 3.5% to 7.5% of the standard deviation of the deviation in profit.

Second, we turn to agency issues faced by loan officers, which might arise from personal connections, poor monitoring of customer officers, or incentives to meet key performance indicators (KPIs). We use hometown ties to capture personal connections between loan officers and borrowers. Personal connections might induce loan officers to be more lenient in loan decisions. To capture the strength of the monitoring of customer officers, who might have incentives to misrepresent loan applications, we identify the loans with customer officers residing in remote areas, where the monitoring cost (e.g., cost of periodic visits from loan officers to the areas) is higher. To capture loan officers' incentives to meet KPI targets, we use the peer pressure loan officers are under—when they lag behind colleagues' performance in the number of approved loans. Based on these proxies, we find that the existence of agency issues is negatively correlated with the value contributed by loan officers. The effect ranges from 9.7% to 16.2% of the standard deviation of the deviation in profit.

Lastly, we consider the behavioral bias of loan officers, including algorithm aversion and mental fatigue. Recent studies suggest that employees do not trust AI recommendations because employees believe their decisions are better. Because AI is better at processing complicated cases and hard information, we use the loans that are more complex—borrowers with loans from other financial institutions—and the loans when loan officers request hard information about borrowers that AI has already processed, to capture loan officers' algorithm aversion. In addition, psychology theory suggests that humans become less effective in decision making when they are subject to mental fatigue. We use the peak hours of loan application processing and the working days around national holidays to proxy for mental fatigue. We find that loan officers' contribution to loan profit is lower when officers have algorithm aversion or experience mental fatigue. The effect ranges from 3.8% to 23.8% of the standard deviation of the deviation in profit.

On July 1, 2024, the focal company implemented a new performance evaluation policy for loan officers by explicitly considering the quality of approved loans, such as the default rate of loans approved in the previous 12 months. Up to this point, the consideration of loan quality in loan officers' performance evaluations is not explicit. We find that after the implementation of the new policy, the use of soft information better enhances the value created by loan officers, and agency issues reduce the value created by loan officers to a less extent. In contrast, after the implementation of the new policy, behavioral bias has a more negative effect on the value created by loan officers, suggesting that the new policy may have unintentionally exacerbated algorithm aversion and mental fatigue for loan officers.

For robustness checks, we conduct several sensitivity tests. First, in the main analysis, we examine the impact of soft information, agency issues, and behavioral bias separately. When we include all the proxies in the regression at the same time, the inferences remain similar. Second, we use alternative approaches to calculate the hypothetical loan profit based on AI recommendations, by changing the hypothetical interest rate, by changing the invoice amount under AI recommendations, or by restricting the sample to loans with a low likelihood of default. We obtain similar inferences. Lastly, instead of calculating the difference in loan profit between human decisions and AI recommendations, we regress loan profit based on human decisions on the hypothetical loan profit based on AI recommendations and the proxies for soft information, agency issues, and behavioral bias. Again, we obtain similar inferences.

Our paper contributes to the literature in three ways. First, it contributes to the literature on AI by examining the value of humans in a real-world AI + Man decision model. We find that humans can contribute by utilizing soft information that is difficult to codify. At the same time, agency issues and behavioral bias can reduce humans' contribution. Given that an increasing number of corporations are investing in AI to assist their employees, our results can thus inform those corporations in better designing the interactions between AI and employees to leverage both humans' and AI's advantage.

Second, the paper contributes to the literature on FinTech and banks. How to design a hybrid model by combining both the recommendations from AI and the inputs from human loan officers are increasingly important questions. Our results can inform FinTech companies and banks by identifying the conditions under which loan officers can add value over an AI system. The results also shed light on how incentive schemes for loan officers can affect officers' incremental contribution over AI recommendations.

Third, we contribute to the literature on the role of soft information in lending decisions by identifying the specific soft information used by loan officers. While the importance of soft information in lending decisions is well-recognized in the literature, prior studies primarily use geographic distance between borrowers and lenders, hierarchical distance between loan officers and supervisors, or relationship lending to capture the *existence* of soft information (Petersen and Rajan 1994, 2002; Stien 2002; Liberti and Mian 2009), due to lack of data on loan decision processes. An exception is Campbell et al. (2019). They identify soft-information-related keywords – those related to borrowers' social, professional, educational, and personal background or those reflecting loan officers' assessments – from loan officers' internal reports and use the ratio of soft-information-related keywords to the total number of words in the reports to capture the extent soft information is used by loan officers. Unlike Campbell et al. (2019), we use loan officers' notes to identify the specific soft information used during loan decisions, such as support from family members, referral by acquaintances, discussions of business prospects, and explanations of financial disputes. Our results indicate that these specific soft information used by loan officers increases the value contributed by officers over AI recommendations.

Our paper is closely related to three papers: Liu (2022), Jansen, Nguyen, and Shams

(2025), and Costello, Down, and Mehta (2020). Liu (2022) and Jansen et al. (2025) compare the performance of AI and human (i.e., AI vs. Man) and find that AI outperforms man. Utilizing data from a small business lender in China, Liu (2022) shows that a machine learning model outperforms loan officers because of humans' constraints in processing hard information. He finds that humans use fewer pieces of hard information, tend to use linear functions, and tend to put a greater emphasis on salient signals. Using a car loan setting, Jansen et al. (2025) find that machine underwriters outperform human underwriters because of agency conflicts faced by human underwriters—their incentives to underprice loans to win the bid in auto loan applications—and human underwriters' poorer ability to process hard information in complex cases. Different from Liu (2022) and Jansen et al. (2025), our paper examines the value of humans on top of AI recommendations in an AI + Man setting. We find that humans' ability to collect and process soft information enhances their contribution and agency issues and behavioral bias, including algorithm aversion and mental fatigue, reduce their contribution. Our analyses thus directly speak to humans' value contribution beyond AI in an AI-assisted decision setting.

Like our paper, Costello et al. (2020) examine an AI + Man setting. They conduct a field experiment in the setting of trade credit extension. While all lenders make decisions based on the recommendations of an automated system provided by a third-party platform, Costello et al. allow half of the trade credit lenders additional discretion in deviating from the machine's recommendations. They find that on average, allowing humans more discretion is associated with a greater reduction in credit risk, especially for private clients without social media accounts. Our paper extends Costello et al. (2020) in several ways. First, we examine whether human intervention in an AI + Man model adds value in loan decisions. Costello et al. focus on trade credit extension. Unlike loans, trade credit extension is tied to the exchange of underlying goods or services. In addition, trade credit has very short duration and does not carry an explicit interest rate. Thus, the results based on the trade credit setting might not be generalizable to loans. Second and more importantly, while Costello et al. primarily focus on *whether* allowing human users a larger discretion from machine-based recommendations affects lending outcomes, we focus on *how* human loan officers add or destroy value. We document that soft information used by, agency issues faced by, and behavioral bias of loan officers affect their value contribution. Thus, our study sheds light on the conditions under which loan officers add or destroy value beyond AI.

2 Institutional background, related literature, and hypothesis development

2.1 Institutional background

The focal company of this study is a licensed lending company headquartered in Zhejiang Province, China. Established in December 2015, it is one of the first small business lending companies approved in the province. Since its inception, the company has positioned itself at the forefront of China's rapidly evolving small business lending sector. Its main clients are small business owners who require several hundred thousand RMB per loan application.

The company's loan approval process has evolved through three distinct stages. In the Expert Decision Phase (2016–2020), credit assessments were carried out by loan officers, who evaluated borrowers based on hard financial data and soft information, such as business reputation and relationship-based insights, which is difficult to quantify but essential for assessing borrowers' creditworthiness in opaque markets.

In the AI Decision Phase (2021–mid 2023), the company transitioned to a fully automated lending model. By leveraging big data and machine learning techniques, it established a risk control system that integrated data on clients' borrowing histories, income patterns, and consumption behaviors. An AI system evaluated all loan applications without human intervention, aligning with the broad FinTech trend toward algorithmic credit scoring and automated decision-making.

Since mid-2023, the company entered the AI + Man Decision Phase, prompted by regulatory and operational considerations. The cap on interest rates newly imposed by Chinese regulators (24% for the company) has limited the AI model's ability to offset high risks with interest rate adjustments. In addition, the competition in the lending market for small businesses has intensified, pushing the company to pursue growth opportunities from a broader set of small business clients, for whom soft information is presumably more important. To address these issues, the company reintroduced expert judgment alongside algorithmic evaluation. Under this hybrid model, loan officers make the final loan approval decisions based on the recommendations of the AI system, whose objective is to assess risk and maximize expected profits.[5] The underlying infrastructure of the AI system is built on state-of-the-art machine-learning methods, including gradient boosting algorithms such as XGBoost. The AI system also includes modules for fraud detection and information verification. The company's proprietary modeling technology has been granted multiple patents in China.

The workflow under the AI+Man model is as follows. An applicant initiates a loan application through the company's mobile app, website, or an offline branch. One of the company's front-end customer officers assists the applicant in completing the application and collects relevant information, including demographic characteristics, socioeconomic attributes, and financial statements of the applicant's business. These inputs are then standardized and submitted to the AI system. The AI model assesses the credit risk using large-scale algorithms that analyze the submitted information as well as over 20,000 variables it collects from various sources.[6] The AI model is trained to assess the risk of the clients and

[5] The loan officers hired in this phase are experienced loan officers from other financial institutions and are not the same as those employed in the Expert Decision phase.
[6] The AI model draws upon an extensive range of data sources, including the People's Bank of China credit registry, data purchased from third-party vendors, and records of individual consumption.

maximize the expected profit of the loan, subject to interest cap. When the borrower's risk is assessed to be higher than acceptable based on the company's loan policies, the AI model will recommend a rejection. The AI model will generate a due-diligence report, which contains a recommendation (approval or denial) and the recommended loan amount, interest rate, and term.[7]

After receiving the AI report, the loan application is randomly assigned to one of the loan officers who conducts an online meeting with both the applicant and the customer officer who handles the application. The meeting typically lasts about 30 minutes, during which the loan officer gathers additional information about the applicant and seeks clarification on some issues when necessary. The loan officer then makes the final decision on whether to approve the loan and for approved loans, the approved credit amount, interest rate, and maturity term, and files the decision report in the system.[8] In the report, the loan officer must document the rationale for the decision and discuss the additional information collected during the meeting.[9] Note that loan officers approve a line of credit to the clients with a term. The client then decides how much and when to draw from the line of credit during the term. The company refers to each drawing as an invoice. On average, clients draw on the credit four times during the approved loan term. We refer to the line of credit as loan throughout the paper for brevity.

After a loan is approved, the focal company's loan service department takes over and sets the repayment schedule for each drawn invoice. If a borrower fails to pay interest or principal in a timely manner, the company will start its debt collection procedure. The

[7] Even when the AI system recommends a rejection, it still provides recommended credit amount, interest rate, and term for the loan application to maximize expected profit subject to interest cap. This information is provided to guide humans' decisions because loan officers make the final loan decisions.

[8] Our setting is similar to Liu's (2022) setting in that the loan amount is typically small. However, unlike Liu's setting, where loan officers' only major decision is the loan amount because maturity term and interest rate are either fixed or determined by the management team, loan officers in our setting need to decide on the loan amount, interest rate, and maturity term.

[9] Human loan officers may also consider the competition for borrowers from other financial institutions, which can result in a lower average interest rate approved by loan officers than AI.

company classifies a loan as in default if the borrower delays the payment of any scheduled principal or interest by more than 10 days. The initial stage of the procedure is by the customer officer who handles the loan application. At this stage, the customer officer contacts the borrower to negotiate a repayment schedule. If this step is unsuccessful, the case is transferred to the company's debt collection department, which undertakes formal debt recovery actions. One of the company's debt collection officers will contact the borrower to request repayment. If this fails, the debt collection team will visit the borrower to ask for repayment. If the borrower remains uncontactable or unwilling to pay, the debt collection department may adopt commonly used public shaming approaches in China, such as contacting the borrower's relatives, acquaintances, or business associates to impose social pressure on the borrower and induce him/her to pay (Bu and Liao 2022; Liao et al. 2025). If all the above efforts fail, the company will sue the borrower and seek the court's help to force the borrower to pay. Because of these strategies and the fact that China does not allow personal bankruptcy, the recovery of at least principals is generally successful.[10]

This AI + Man decision model provides an ideal setting to investigate the value of human intervention in lending decisions and examine how soft information, agency issues, and behavioral frictions shape the value of human intervention.

2.2 Literature review

A growing literature examines the economic consequences of AI in financial decision-making, paying particular attention to how AI complements or substitutes for human judgment. For example, Grennan and Michaely (2021) find that AI analysis adopted by FinTech firms can improve stock price informativeness and can potentially substitute for traditional sell-side research. Coleman, Merkley, and Pacelli (2022) document that the

[10] Because of the risk control policy and the effectiveness of debt collection procedure of the focal company, the number of actual defaults (i.e., non-collection of principals) in our sample is very small. Thus we are unable to examine how human intervention affects the likelihood of default. Instead we focus on the variation of loan profitability to capture the value of human involvement.

recommendations of Robo-Analysts—human analyst-assisted computer programs—contain positive long-term investment value and are less biased and more frequently updated than human analysts' recommendations. van Binsbergen et al. (2023) use their AI models to detect biases in human analysts' forecasts. Exploiting the one-month ban on ChatGPT in Italy, Bertomeu et al. (2025) find that analysts integrate large language models (LLMs) into their workflow, which influences forecasting behavior and outcomes. This line of research supports the notion that AI is useful for security analyses. Extending to firm decision setting, prior research documents productivity gains from AI adoption (e.g., Fedyk et al. 2022; Jansen et al. 2025). For example, Babina et al. (2024) find that firms' investments in AI can increase sales growth and market valuation, partially because of the positive effect of AI investments on product innovation.

Another line of literature investigates how AI shapes the labor market. For example, Abis and Veldkamp (2024) find that firms adopting AI technologies rely less on labor in their production function. Using establishment-level data, Acemoglu et al. (2022) document that as establishments increase investments in AI, the hiring of non-AI positions decreases.

Particularly relevant to our study, Cao et al. (2024) compare the performance of human analysts with an AI analyst built by the authors (i.e., AI vs. man). They further compare the performance of the AI analyst with that of a "Man + AI" analyst, which uses human analysts' forecasts as an additional input to the AI analyst. They find that human inputs are more valuable when covering firms that are more illiquid, have more intangible assets, and have higher earnings volatility and distress risk. Complementing Cao et al., we focus on the value of human intervention in an AI + Man setting, where human loan officers make the loan decision based on the recommendations of an AI system. That is, unlike the setting in Cao et al., humans make the final decisions in our setting, which is close to the scenarios where companies provide AI assistance (i.e., "robot assistants") to their employees in practice as the focal company does. Our setting allows us to examine the value of humans over an AI-based decision model and how the value of humans varies with the soft information used by, agency costs faced by, and behavior bias of, human loan officers. In addition, our evidence speaks to the synergies of man and machine in the loan decision setting.

2.3 Hypothesis development

A central question in the adoption of an AI + Man lending decision model is whether human intervention improves or impairs the value of loans. Theoretical and empirical research provides reasons for both positive and negative effects, leading to a prediction of an ambiguous net effect. As a result, our focus in this paper is not on the average effect, but on the conditions under which human intervention is more or less likely to improve on AI's recommendations. We focus on three dimensions: (1) the use of soft information, (2) the agency conflicts faced by loan officers, and (3) behavioral bias of loan officers including algorithm aversion and mental fatigue.

The benefits of soft information. Conceptually, because AI is a prediction technology, humans can contribute by making judgements based on inputs that cannot be easily quantified or codified (Agrawal et al. 2019). While modern machine-learning algorithms excel at processing hard information—verifiable, structured data such as income, education, and spending patterns, they currently have limited capabilities to acquire and interpret soft information, which often requires personal interactions, contextual judgment, and tacit knowledge. Soft information—qualitative, non-standardized, or relationship-based information—is critical for assessing borrower risk (Liberti and Petersen 2019). The value of soft information is well documented in the literature (e.g., Campbell et al. 2019). Liberti and Petersen (2019) highlight that soft information plays a critical role when the borrower's quality cannot be fully inferred from historical or quantitative indicators, especially for small-business and relationship lending. Gerken and Painter (2023) show that local or contextual knowledge can improve financial decision-making by enriching the information set beyond what standardized analytical tools capture. These findings suggest that human loan officers can contribute complementary insights over AI systems through the use of soft information.

In the small business lending setting, soft information can include impressions of borrowers' integrity, the credibility of their narratives, and nuanced aspects of business conditions, all of which are difficult to codify or quantify algorithmically at the current times. For example, in our setting, loan officers meet with loan applicants, and loan decisions can be influenced by loan officers' perceptions of the applicants and the additional information acquired during the meeting, such as family support and contextual business information. In such cases, the marginal value of human intervention is likely to be higher.

In sum, loan officers can enhance the value of loans by collecting and using soft information that complements the AI model. Our first hypothesis is thus stated as follows:

H1: Ceteris paribus, human intervention is more likely to improve the value of loans when soft information is used.

Note that H1 might not hold for two reasons. First, while imperfect predictors, AI models are increasingly sophisticated. It might be able to incorporate some factors that are correlated with soft information or quantify soft information (Liberti and Petersen 2009; Sutherland 2020), reducing the value of soft information collected by loan officers. Second, AI models can collect information from an increasing number of sources. The increasing amount of information used by AI models can reduce the marginal value of soft information collected by loan officers.

Despite the potential benefits of human intervention, theoretical and empirical studies suggest that human intervention can also reduce lending decision quality. We examine two possibilities, agency problems and behavior bias.

Agency problems faced by loan officers. Loan officers may face misaligned incentives that distort assessments (e.g., Liberti and Petersen 2019). These frictions reduce the quality of human adjustments to AI recommendations, leading to poorer loan decisions. For example, personal connections between loan officers and borrowers, such as shared community ties, may create favoritism, reducing the value of human intervention. For another example, loan officers may lack incentives to monitor customer officers who can misrepresent information or hide unfavorable information of borrowers. In addition, loan officers usually have loan approval targets, creating incentives to approve marginal borrowers to meet targets. In these cases, human intervention can reduce lending decision quality relative to AI recommendations. Accordingly, when agency problems faced by loan officers are more severe, human adjustments to AI recommendations are less likely to add value. Our second hypothesis is stated as follows:

H2: Ceteris paribus, human intervention is less likely to increase the value of loans when loan officers face more severe agency problems.

Behavioral bias. Behavioral bias related to algorithm aversion and mental fatigue can also reduce the value of human intervention. Algorithm aversion refers to the phenomenon that workers avoid or reduce the reliance on AI assistance even when doing so reduces task performance (e.g., Dietvorst et al. 2015; Commerford et al. 2022; Commerford et al. 2024; Bockstedt and Buckman 2025). Jussupow, Benbasat, and Heinzl (2024) show that individuals frequently underweight or reject algorithmic inputs—even when the algorithm is demonstrated to be more accurate. This phenomenon exists our setting. For example, one of the loan officers shared that "with more than ten years of experience, we have better judgement than the model." It is also possible that human decision-makers revise algorithmic recommendations mainly to demonstrate their value and to keep their jobs. In the lending context, algorithm aversion may lead loan officers to revise AI recommendations excessively, thereby reducing the value of human intervention.

Research in behavioral economics and accounting demonstrates that mental fatigue can reduce the quality of professional judgment, particularly under time pressure or with high workloads (e.g., Baumeister et al. 1998; Hirshleifer et al. 2019). Humans with mental fatigue tend to rely on heuristic decision-making process, that is, making decisions based on intuition, or easy and quick cognitive processes, instead of slow, rigorous reasoning processes (Stanovich and West 2000). For example, Hirshleifer et al. (2019) document that financial analysts become less accurate in their earnings forecasts when their workload accumulates through the day, suggesting that human judgment worsens under cognitive overload. In the lending context, loan officers who have peak-period workloads may experience fatigue that reduce their ability to effectively evaluate borrowers or critically assess AI's recommendations. With mental fatigue, loan officers may rely on heuristics, become less attentive to risk indicators, or rely on AI recommendations without appropriate scrutiny. Our last hypothesis is thus stated as follows:

H3: Ceteris paribus, human intervention is less likely to increase the value of loans when loan officers are more likely to be subject to behavioral bias.

In sum, loan officers may complement AI by collecting and using soft information and contextual insights, yet they may also hinder loan performance because of misaligned incentives, ineffective use of algorithmic recommendations, or mental fatigue. While the net effect of human intervention in AI-assisted lending decisions is ambiguous, an understanding of the conditions under which loan officers contribute or destroy value can inform governments and companies in an increasingly integrated AI-assisted economy.

3 Sample and research design

3.1 Sample

We start with the population of 28,198 loans approved by the focal company in 2024. We do not observe loan applications that were not approved because under the government regulation, the company is not allowed to keep the information about rejected loan applications due to data privacy concern.[11] We do not consider loans approved in the second half of 2023 because loan officers likely need time to learn how to work with AI recommendations. We do not consider loans approved after 2024 either because we need time to evaluate the outcomes of the approved loans. We drop loans with invoices (e.g., actual drawings from the credit line by borrowers) that have not matured at the time of data collection. Note that while the approved loan credit might have a one-year or two-year term, invoices typically have a shorter term. We further drop loans with missing variables, loans approved by inexperienced loan officers, i.e., those who approved fewer than 100 loans in 2024,[12] and singletons due to the inclusion of fixed effects. The final sample includes 18,490 loans approved by 24 loan officers.[13] Table 1 provides the summary of sample selection process.

Panel A of Table 2 provides descriptive statistics on loan outcomes based on AI recommendations and loan officers' decisions. The average loan amount is CNY126,947 under AI recommendations, lower than the average based on loan officers' decisions (CNY 161,846). The average interest rate recommended by AI is 18.15%, higher than that decided by loan officers (14.51%). The average term of the line of credit recommended by AI is 16.58 months, shorter than that approved by loan officers (24.62 months). All the differences are significant at the 1% level. These statistics indicate that on average, loan officers approve a larger amount of loans with a lower interest rate and a longer term than AI recommendations.

[11] See Article 47 of the *Personal Information Protection Law of the People's Republic of China* (2021). Based on our discussions with the focal company, for loans approved by AI, loan officers may adjust loan parameters but rarely reject the loans. For loan rejected by AI, loan officers may approve the loan and adjust loan parameters, or loan officers may also reject the loan. Thus, loan applications not approved (and not captured by the sample) consist of loans rejected by both AI and humans. Since AI and humans agree in these omitted cases, these omissions will not affect our inferences regarding the incremental value of humans over AI.

[12] The average (median) number of approved loans across loan officers in 2024 is 792 (720).

[13] Loan applications are randomly assigned to loan officers. Separately, because loan officers are relatively homogeneous in education level and work experience, we are unable to conduct cross-sectional analyses based on loan officers' characteristics.

3.2 *Loan profit calculation*

Following Jansen et al. (2025), we calculate the profit of each invoice of the loan using realized or expected cash flows as follows:

$$Profit = \sum_m [PV(Payment_m) + PV(Recovery|Default)] - Initial\ Investment, \quad (1)$$

where *Profit* is the profit of the invoice for the company, *m* denotes month m, *PV* is the present value calculated using the corresponding Shanghai Interbank Offered Rate (Shibor), $Payment_m$ is the payment made by the borrower in month m, which includes both the interest payment and the repaid principal of the loan, *Recovery|Default* is the recovered principal and accrued interest upon default, and *Initial Investment* is the dispensed invoice amount.

To evaluate the deviation of loan officers' decisions from AI's recommendations, we calculate *Profit* based on both loan officers' decisions and AI's recommendations. We calculate *Profit* based on loan officers' decision using the dispensed invoice amount and realized cash flows and denote it as $Profit_{-Man}$. Appendix A provides an example of the calculation.

The calculation of loan profit based on AI's recommendations requires additional assumptions. First, we assume that the hypothetical invoice amount under AI's recommendations is proportional to the actual invoice amount, with the proportion calculated as AI's recommended credit amount over the credit amount approved by the loan officer. The premise underlying this assumption is that the amount borrowed increases with the size of the credit limit. We validate this assumption by regressing the total invoice amount under each loan on the approved credit amount and borrower characteristics. We find that the coefficient on the credit amount is 0.901 with a *t*-statistic of 72.54. Second, the repaid principal under AI's recommendation is similarly calculated in proportion to the actual repaid principal. The assumption is, had the company followed AI's recommendation, the likelihood of the borrower's default would be the same as the observed default likelihood under loan officers' decisions. As discussed later, our inferences remain the same if we focus on the sample where default is not a significant concern. Also note that the repayment schedule is set independently by the loan service department after a loan is approved. Third, we use AI's recommended interest rate for loans approved by AI. For loan applications rejected by AI, we assume that the firm accepts AI's decision, and the funds allocated for these loans will be reallocated to other borrowers. Thus, for loan applications rejected by AI, we use the average interest rate charged on loans approved in the previous month. As reported in Panel A of Table 2, the hypothetical AI interest rate is 16.39%, still higher than the rate set by loan officers.

Based on these assumptions, we derive the hypothetical invoice amount, repaid principal, paid interest, and calculate *Profit* accordingly, denoted as $Profit\text{-}_{AI}$. Please refer to Appendix A for an example of the calculation of $Profit\text{-}_{AI}$.

Because we do not observe the loan outcomes under AI's recommendations, the estimation of $Profit\text{-}_{AI}$ contains measurement errors. However, the errors are unlikely to introduce systematic bias to the results for H1-H3 because they are unlikely to be correlated with the proxies for soft information, agency issues, or behavioral bias. In addition, we conduct sensitivity tests by using alternative measurements and obtain the same inferences, as discussed in Section 5.

Lastly, we calculate the deviation in profit of loan officers' decisions from AI's recommendations as follows:

$$Deviation\_Profit\ (\%) = 100 \times \frac{Profit_{-Man} - Profit_{-AI}}{Invoice\ Amount_{-Man}}, \tag{2}$$

where $Invoice\ Amount_{\_Man}$ is the total dispensed loan amount to the borrower. For loans with multiple invoices (i.e., multiple drawings by borrowers), we aggregate the profit at the invoice level to obtain the loan-level profit. *Deviation_Profit* (%) captures the value created by human loan officers in the AI-assisted lending decisions.

3.3 Research design

To test hypotheses H1-H3, we estimate the following regression:

$$\begin{aligned} Deviation\_Profit(\%)_l \\ = \alpha + \beta_1 Soft\_Information_l + \beta_2 Agency\_Issue_l \\ + \beta_3 Behavioral\_Bias_l + \varepsilon_l \, , \end{aligned} \tag{3}$$

where subscript *l* represents loan *l*. *Soft_Information* is the proxy for soft information collected and used by loan officers in the approval process. *Agency_Issue* is the proxy for agency issues faced by local officers. *Behavioral_Bias* is the proxy for behavioral bias (i.e., algorithm aversion and mental fatigue) of loan officers. We describe these proxies in the next section when we report the analyses. A positive coefficient on these variables indicates that these variables enhance the value contributed by human loan officers, while a negative coefficient indicates that they reduce the value contributed by human loan officers.

Because *Deviation_Profit* is calculated based on the difference between the loan officer's decision and AI recommendation for the same loan, loan characteristics and borrower characteristics are held constant. As a result, even if some of the proxies for soft information, agency issues, or behavior bias might reflect certain loan characteristics such as risk, our inferences are unlikely to be driven by the alternative explanation related to loan characteristics. We include month, borrowers' industry, and loan officer fixed effects to control for time trend, time-invariant industry characteristics, and innate attributes of loan officers. The *t*-statistics are calculated using standard errors clustered at the loan officer level. The inferences remain the same if we use wild cluster bootstrap p-values.

4 Main tests

4.1 Value of Man in AI + Man

We start the analyses by examining the average value contributed by human loan officers over AI recommendations. Specifically, we compare the profit of loans based on AI recommendations and loan officers' decisions. Panel A of Table 3 reports the results. We find that the average profit based on AI recommendations is CNY 12,109, and the average profit based on loan officers' decisions is CNY 13,153. The difference is significant at the 1% level. These statistics suggest that on average, human loan officers contribute positively to loan decisions.

We obtain the same conclusion based on the deviation in profit of loan officers' decisions from AI recommendations, *Deviation_Profit (%)*. We winsorize *Deviation_Profit (%)* at the 1% and 99% percentiles. As reported in Panel B of Table 3, the mean of *Deviation_Profit (%)* is 0.209. As shown in Figure 1, while the distribution of *Deviation_Profit (%)* centers around zero, there is a long tail on the left with extreme negative values. After winsorization, the minimum of *Deviation_Profit (%)* is -23.59, while the maximum is 11.40 (untabulated). In addition, *Deviation_Profit* (%) has a large variation: the standard deviation is 6.164, and for about 43% of the loans, human intervention reduces the value of loans. These statistics suggest that it is important to understand the conditions under which human loan officers add or reduce value beyond AI recommendations.

To understand what drives the difference in profit, we run a regression of *Deviation_Profit (%)* on the difference in loan parameters between loan officers and AI — amount (*Deviation_Amount*), interest rate (*Deviation_Rate*), term (*Deviation_Term*)—and the likelihood of loan officers overriding AI's recommendations of loan denial (*Override*). Panel B of Table 3 reports the descriptive statistics on these variables. The mean of *Deviation_Amount* is 2.134 times AI's recommended loan amount, the mean of *Deviation_Rate (%)* is -1.875,[14] and the mean of *Deviation_Term* is 8.043 months. The mean of *Override* is 0.609.[15]

[14] In the following analyses, unless noted, we use the hypothetic AI rate, which is set as the AI recommended interest rate for loans approved by AI and the average interest rate of the loans approved in the previous month for loans rejected by AI.
[15] An untabulated analysis indicates that loans denied by AI but approved by loan officers have a smaller amount, lower interest rate, longer term, and lower profit than those approved by both AI and loan officers.

Panel C of Table 3 reports the regression results. We control for month, industry, and loan officer fixed effects. When the four variables (*Deviation_Amount*, *Deviation_Rate*, *Deviation_Term*, and *Override*) are included individually in Columns (1)-(4), all of them are significantly positively correlated with *Deviation_Profit (%)* (t = 63.13, 20.14, 5.82, and 14.60, respectively). These results suggest that when loan officers increase loan amount, interest rate, and term and override AI's denial decisions, their contribution to loan profit is higher. When we include all the four variables at the same time, as reported in Column (5), the inferences remain the same except that *Override* is no longer significant.

In sum, the results suggest that on average human loan officers contribute positively to loan decisions. Moreover, the deviation in profit has a large variation. Thus, it's important to understand the conditions under which humans add or subtract value by deviating from AI's recommendation.

4.2 Soft information and the deviation in profit – Tests of H1

H1 states that the value of human loan officers in an AI + Man decision-making setting increases with the use of soft information by loan officers. While the concept is straightforward, the challenge in the literature has been how to capture soft information used by loan officers. We overcome this challenge by obtaining access to the comments made by loan officers when approving loans. Based on these comments, we construct five proxies for soft information: (1) an indicator for the length of the comments (*Comment_Long*); (2) the discussion of the prospect of borrowers' businesses (*Business_Prospect*); (3) the support for borrowers from family members (*Family_Support*); (4) the discussion of the impact of lawsuit faced by the borrowers or their cash flow anomalies on their creditworthiness (*Credit_Anomalies*);[16] and (5) the referral from existing clients or loan officers'

[16] The lawsuits are usually related to financial disputes on payment arrangement of business transactions, guarantees, or contract enforcement. They are usually not related to borrowers' financial or criminal misconduct.

acquaintances (*Referral*).[17] Appendix B provides detailed definitions of these variables.

First, given that the comments are mainly about information not used by AI, the longer the comments are, the more soft information loan officers collect and use in their decisions. We classify the comments with length (i.e., the number of Chinese characters) in the top 25% of the sample distribution as long comments. Second, some comments discuss the prospect of the borrowers' businesses, such as the borrowers' expertise, government policy support, and growth potential. Such information can help loan officers to better evaluate borrowers' creditworthiness. For example, one loan officer notes that "The borrower demonstrates strong operational capability, exhibits favorable growth prospects, with the business currently in an expansion phase." Another notes that "The borrower has extensive experience in the core business, maintains a stable customer base, and is able to expand into new markets by training apprentices." Third, the support from family members can increase borrowers' ability to pay interests and principal. The old Chinese saying that "Sons will pay the father's debt" highlights that the support from family members is an important form of loan guarantee. For example, one loan officer notes that "The borrower's son demonstrates strong support for the family business and may contribute to repayment if necessary." Another notes that "The client's two daughters have low debt and are willing to support the business." At the same time, such support is not guaranteed, and loan officers need to evaluate how credible the family support is in each case. Fourth, some comments include loan officers' evaluation on the impact of borrowers' legal disputes or cash flow anomalies on their creditworthiness. For example, one loan officer notes that "The borrower expects the legal dispute to be resolved within one to two months, after which the business will start generating cash flow." Another notes that "The increase in liabilities in the last two years is primarily

[17] To codify soft information, we first randomly select 2,000 comments. We then assign three research assistants to identify topics within the comments independently and cross-validate the classifications for consistency and accuracy. Lastly, we use the classifications of topics in these 2,000 comments as the training dataset to fine-tune a BERT model, which is then used to identify topics for the remaining comments.

attributable to the renovation of the homestead, which has now been completed." Lastly, in China, referring a friend to a business has the implicit "guarantee" of the person's integrity and creditworthiness. For example, one loan officer notes that he "consulted Ms. Pan, a customer officer who previously handled the loan application of the borrower's spouse. Ms. Pan indicated that the couple is generally credible."

Panel B of Table 2 reports descriptive statistics on these variables. About 25.3% of the loans have long comments from loan officers, 14.1% of loans discuss borrowers' business prospects, 12.7% of loans indicate family support, 15.0% of loans provide explanations for financial dispute or cash flow anomalies that might affect borrowers' creditworthiness, and 34.3% of loans are referred by existing clients or acquaintances. Appendix C reports the correlation between these variables. The correlation coefficients are generally small, indicating that they capture different dimensions of soft information.

We then use these five proxies to explain *Deviation_Profit* and report the regression results in Table 4. We find that the coefficients on all five proxies (*Comment_Long*, *Business_Prospect*, *Family_Support*, *Credit_Anomalies*, *Referral*) are significantly positive (t = 2.65, 2.05, 2.92, 2.95, and 2.70, respectively). Compared with other loans, loan officers' contribution in profit is higher by 5.9% (= 0.366/6.164) of the sample standard deviation of *Deviation_Profit (%)* for loans with long comments from loan officers, by 3.5% for loans with information on borrowers' business prospects, by 7.5% for loans indicating support for borrowers from family members, by 7.1% for loans with explanations about borrowers' creditworthiness, and 3.6% for loans referred by existing clients or acquaintances.

In sum, these results indicate that human loan officers can create value in an AI + Man decision model through the use of soft information.

4.3 Agency issues and the deviation in profit – Tests of H2

H2 states that the value of human loan officers in an AI + Man setting decreases with agency issues faced by loan officers. We construct three proxies for agency issues faced by loan officers: (1) the hometown tie between loan officers and the borrowers (*Hometown_Tie*), (2) the location of customer officers in remote areas (*Remote_Area*), and (3) the peer pressure in achieving KPIs (*Peer_Pressure*). Appendix B provides detailed definitions of these variables.

First, hometown tie—the loan officer and the borrower sharing the same place of origin—is one of the most widespread sources of networks in the Chinese society. Prior research finds that hometown ties affect a wide range of decisions, such as politicians' career progression (Fisman et al. 2020), the quality of government monitoring (Chu et al. 2021), and audit quality (Deng, Zhang, and Liu 2023). We thus expect loan officers to be more lenient with borrowers from the same hometown (e.g., by charging a lower interest rate), reducing the value of loans. Second, the focal company has a policy of requiring loan officers to periodically visit and monitor customer officers in order to ensure that customer officers do not misrepresent loan applications, e.g., to help borrowers with a poor credit record to obtain loans. While customer officers reside in the same province as the borrowers, loan officers reside in the headquarters of the company in Zhejiang province. Because traveling to remote areas is time-consuming for loan officers, loan officers might travel less to remote areas and as a result monitor customer officers in those areas less effectively, reducing the quality of loans from such areas.[18] Third, the number of approved loans is a critical KPI for the loan department. Hence, there is pressure for the loan officers as a group to achieve KPIs and for the individual loan officers not to lag in performance. Thus, we expect the loan officers who are lagging in processing loan applications to be under pressure to increase the number of

[18] The focal company regards Yunnan, Guizhou, Gansu, Shaanxi provinces as remote areas. The company does not have businesses in other similarly remote areas such as Xinjiang and Tibet.

loans approvals at the expense of decision quality.[19]

As reported in Panel B of Table 2, 5.0% of the loans are approved by loan officers with hometown ties with borrowers, 12.6% of the loans are processed by customer officers from remote areas, and 41.1% of the loans are approved when loan officers are under peer pressure. As indicated in Appendix C, these variables are not highly correlated with each other, indicating that they capture different dimensions of agency issues faced by loan officers.

Table 5 reports the regression results. We find that the coefficients on *Hometown_Tie*, *Remote_Area*, and *Peer_Pressure* are significantly negative (t = -4.44, -4.70, and -5.98, respectively).[20] Compared with other loans, loan officers' contribution to profit is lower by 16.2% (= -0.998/6.164) of the sample standard deviation for loans approved by loan officers who have hometown ties with borrowers, by 9.7% for loans processed by customer officers in remote areas, and by 11.4% for loans approved by loan officers under peer pressure.

These results indicate that loan officers faced with agency issues, arising from personal connections with borrowers, weak monitoring of customer officers, or peer pressure to achieve KPIs, create less value, or destroy value, by deviating from AI's recommendations.

4.4 Behavioral bias and the deviation in profit – Tests of H3

H3 states that the value of human loan officers in an AI + Man decision setting decreases with behavioral bias of loan officers. We consider two dimensions of behavioral bias: (1) loan officers' algorithm aversion, and (2) loan officers' mental fatigue.

With respect to loan officers' algorithm aversion, we construct two proxies: (1) loans applied by borrowers who have loans from other financial institutions (*Multi_Loan*), and (2)

[19] The focal company intentionally creates peer pressure among loan officers. It installed a large screen in the meeting room to display real-time approval progress, allowing loan officers to observe their colleagues' workloads and progress in real time.

[20] It is possible that hometown ties enable loan officers to collect some soft information about the borrowers. If so, the coefficient on *Hometown_Tie* should be positive, opposite to our finding.

the request of hard information by loan officers (*Infor_Request*). When a borrower has loans from other financial institutions, the loan officer needs to consider the implications of a larger amount of funding for the borrower's business, the seniority of the loan from the focal company, and the implications of having other loans for the borrower's ability to pay interests and principal. In such situations, AI's access to extensive information about borrowers and adaptation at processing complicated cases (Liu 2022) can amplify AI's advantage. Loan officers should be aware of AI's relative advantage in these cases. The focal company's dashboard for loan officers display a special tag for such loans. If loan officers suffer from algorithm aversion, they may make similar adjustments to AI recommendations as for other loans or even make greater adjustments than for other loans, likely leading to value-destroying decisions.[21] Similarly, the focal company's AI system processes over 20,000 pieces of information. When loan officers request to review information that has been processed by the AI system (e.g., the business permit and financial statements), the loan officers likely suffer from algorithm aversion, which can reduce the quality of their loan decisions. As reported in Panel B of Table 2, 38.3% of the loans are applied by borrowers who have loans from other financial institutions, and 17.6% of the loans are approved by loan officers who have requested information already processed by the AI system. As reported in Appendix C, the two variables are positively correlated with each other, but the correlation coefficient is very small (0.036).

Columns (1) and (2) of Table 6 present the regression results. The coefficients on *Multi_Loan* and *Infor_Request* are significantly negative (t = -12.07 and -3.52, respectively). Compared with other loans, loan officers' contribution to profit is lower by 23.8% (= -1.465/6.164) of the sample standard deviation for loans where the borrowers have loans from

[21] Note that while the loans to borrowers who also have loans from other financial institutions are potentially risker than other loans, the calculation of *Deviation_Profit* controls for the impact of loan characteristics, because *Deviation_Profit* captures the incremental value contributed by loan officers over AI recommendations. This point also applies to the other proxies that might be correlated with loan characteristics.

other institutions and by 7.8% for loans approved by loan officers who have requested information processed by the AI system.

With respect to mental fatigue, we construct two proxies: (1) the hours during which loan officers approve the highest number of loans, i.e., 10-11am and 4-5pm (*Peak_Hour*), and (2) the working days around holidays (*Around_Holiday*). Psychology theory suggests that mental fatigue occurs after a prolonged period of cognitive activities and can impair humans' cognitive ability and the quality of decision-making (e.g., Baumeister et al. 1998; Shleifer et al. 2019). Thus, loan officers will be more likely to experience mental fatigue and make poor decisions when processing a large number of loan applications in a short period or around holidays. As reported in Panel B of Table 2, 25.6% of the loans are approved during peak hours,[22] and 5.0% of the loans are approved around holidays. These two proxies are not significantly correlated with each other, as reported in Appendix C.

Columns (3) and (4) of Table 6 present the regression results based on these two proxies. The coefficients on *Peak_Hour* and *Around_Holiday* are significantly negative (t = -2.68 and -2.60, respectively). Compared with other loans, loan officers' contribution to profit is lower by 3.8% (= -0.233/6.164) of the sample standard deviation for loans approved during peak hours and by 6.2% for loans approved around holidays.

Overall, these results suggest that loan officers' contribution over AI's recommendations is lower when they have algorithm aversion or experience mental fatigue.

4.5 Separate analyses for loans approved in the first and second half of 2024

The focal company did not have a formal performance evaluation policy for loan officers when it first transitioned to the AI + Man decision model. Loan officers were compensated with a fixed salary individually negotiated with the company and a bonus.

[22] In the focal company, loan officers start working at 8am. The official working hours end at 5pm. However, the company provide free dinner at 5pm. Many loan officers stay for dinner and work afterwards for a couple of hours. Thus, the number of working hours is longer than 8.

While the loan officers, who are experienced hires from other financial institutions, understand the importance of the quality of loans, the quality of loans is not explicitly considered in their performance evaluation, potentially leading loan officers to focus more on the number of loan approvals.[23] On July 1, 2024, the focal company started to implement a new policy on the evaluation of loan officers. The evaluation considers both the number of loans approved and the default rate of loans approved in the previous 12 months, with both dimensions receiving equal weights. That is, the company explicitly considers both the quantity and quality of approved loans for the second half of 2024, while up until this point the consideration of loan quality is not explicit. This new policy encourages loan officers to pay closer attention to the quality of loan decisions. Increasing the number of approved loans at the expense of loan quality will directly reduce their compensation in the future. We thus expect that loan officers are more diligent in approving loan applications in the second half of 2024. The impact of soft information, agency issues, and behavioral bias might be different for the second half of 2024 compared with the first half. To test this, we replicate the main tests separately for loans approved in the first half of 2024 and in the second half of 2024.

Table 7 reports the regression results. Panel A focuses on soft information. We find that soft information is significantly more positive in explaining *Deviation_Profit* in the second half than in the first half; the differences in the coefficients on *Business_Prospect* and *Family_Support* are significantly different from zero. These results suggest that soft information, particularly the information about borrowers' business prospects and family support for loan payment, plays a more important role in loan officers' contribution to loan profits when loan quality is more important to loan officers.

Panel B reports the regression results for agency issues. We find that agency issue

[23] This focus can be strengthened by the fact that up until the first half of 2024, the loan department and the loan origination department are combined and managed by the same executive. For the loan origination department whose main task is to solicit loan applications (by customer officers), the quantity of loans approved is the key performance measure. The two departments are separated approximately in the second half of 2024.

proxies are significantly negatively associated with *Deviation_Profit* for loans approved in the first half of 2024. Only the coefficient on *Peer_Pressure* is significantly negative for the second half of 2024. The coefficients on *Hometown_Tie* and *Remote_Area* are more negative in the first half of 2024 than in the second half of 2024, although that on *Peer_Pressure* is more negative in the second half of 2024. These results suggest that emphasizing loan quality can partially mitigate the agency issues faced by loan officers but exacerbate peer pressure.

Lastly, Panel C reports the regression results for behavior bias. We find that while the two algorithm aversion proxies are important in explaining *Deviation_Profit* in both periods, the coefficients are significantly more negative for the second half of 2024 than for the first half. This result suggests that loan officers rely less on AI recommendations when their compensation is more closely tied to the quality of loan decisions. In addition, mental fatigue proxies are only significant for the second half of the year, suggesting that loan decisions take a greater toll on loan officers when loan quality is an important consideration.

Overall, these results indicate that a well-designed performance evaluation framework can shape loan officers' decisions. The emphasis on loan quality can induce loan officers to rely more on soft information, partially mitigate agency issues faced by loan officers, but exacerbate the negative impact of algorithm aversion and mental fatigue.

5 Additional analyses

In this section, we report a few additional analyses and sensitivity tests to enrich the results and to ensure that the results are robust to alternative research design choices.

5.1 The robustness of the results when all proxies are included at the same time

To investigate whether results on individual variables examined above hold when other variables are also included, we estimate a regression by including all proxies for soft information, agency issues, and behavior bias at the same time. Table 8 reports the regression results. Columns (1), (2), and (3) include all proxies for soft information, agency issues, and behavior bias, respectively. We find that the results continue to hold. When we include all proxies at the same time, as reported in Column (4), we find that the results continue to hold for all proxies except *Business_Prospect* and *Infor_Request*, which are marginally significant (t = 1.56 and -1.30, respectively). For ease of result presentation, we use the model in Column (4) for the remaining sensitivity tests.

5.2 Alternative approaches to calculate profit for loans under AI recommendations

In the main analysis, we make certain assumptions when calculating loan profit under AI recommendations: (1) using the average interest rate charged on the loans approved in the previous month for loan applications rejected by AI, (2) assuming that the default risk under AI recommendations is the same as the actual default risk, and (3) assuming that the invoice amount under AI recommendations is in proportion to the actual invoice amount. In this section, we relax these assumptions and evaluate whether the inferences continue to hold.

Using alternative interest rate for loans rejected by AI. We calculate the hypothetical profit based on AI recommendations using the interest rate recommended by AI for loans approved by AI and using the average interest rate charged on all loans approved in the previous month for the loans rejected by AI. In this section, we replicate the main analyses by using the interest rate recommended by AI for all the loans to calculate *Deviation_Profit*; for loan rejected by AI, had AI recommended loan approval, that would be the interest rate charged on the borrower. We replicate the tests of H1-H3 using the new measure of *Deviation_Profit*. Column (1) of Table 9 reports the regression results. The inferences remain the same, except that the coefficients on *Referral*, *Infor_Request*, and *Around_Holiday* are insignificant.

Restricting to the sample of loans that are unlikely to default. To address the concern that our inferences might be affected by the differential default risk under AI's recommendations (i.e., the hypothetical default risk had the company followed AI's recommendations), we focus on borrowers with a low likelihood of default, for which our assumption of similar default risk under loan officers' decisions and AI recommendation is more likely to hold. For this purpose, we replicate the main analysis using the sample of loans that do not default. The company designates a borrower to be in default if the borrower delays the payment of interest or principal by 10 days or more. Column (2) of Table 9 reports the regression results. The inferences remain the same, except that the coefficients on *Business_Prospect*, *Referral*, and *Infor_Request* are insignificant.

Using an alternative measure of invoice amount under AI recommendations. In this sensitivity test, we assume that the invoice amount under AI recommendations is the same as the actual invoice amount as long as it is lower than the credit limit approved by AI; otherwise we use the credit limit of AI as the invoice amount. The underlying assumption is that a borrower will borrow a fixed amount regardless of the credit limit. We do not believe that this assumption is reasonable as borrowers tend to borrow more when credit limit is higher. Nevertheless, we replicate the main analysis, recalculating the deviation in profit by changing the invoice amount and the payment of interest and principal accordingly. Note that this approach mostly assumes away the impact of the difference in approved credit limit between AI and human loan officers.

Column (3) of Table 9 reports the regression results. We find that the inferences remain the same, although the results are much weaker than those reported earlier. Two soft information proxies, all agency issue proxies, and one behavioral bias proxy remain significant. The weaker results are not surprising given that the most important difference between AI recommendations and loan officers' decisions—the difference in the approved credit limit—is assumed away.

5.3 Alternative research design

Lastly, we use an alternative research design to test H1-H3. Instead of taking the difference between $Profit_{\_Man}$ and $Profit_{\_AI}$, we regress ln($Profit_{\_Man}$) on ln($Profit_{\_AI}$) and the proxies for soft information, agency issues, and behavior bias. In this approach, the measurement error in $Profit_{\_AI}$ should have a smaller impact on the coefficients on the proxies for soft information, agency issues, and behavior bias. As reported in Table 10, our inferences remain the same, although the results are weaker. Four out five soft information proxies, two out of three agency issue proxies, and one out of four behavior bias proxies are significant. The results indicate that after controlling for loan profit under AI recommendations, the loan profit based on loan officers' decisions increase with soft information, decrease with agency issues, and decrease with behavior bias.

6 Conclusion

In this paper, we examine the value of humans in an AI + Man model by exploiting the proprietary data from a FinTech company that uses a hybrid AI + Man model to process loan applications from small businesses. We find that the value of human loan officers increases with the use of soft information, decreases with agency issues, and decreases with loan officers' algorithm aversion and mental fatigue. We further find that a performance evaluation policy emphasizing loan quality can induce loan officers to rely more on soft information, partially mitigates agency issues faced by loan officers, but exacerbates the issues caused by behavioral bias. The results are robust to alternative research design choices. The study contributes to the AI literature by identifying the conditions under which humans add or destroy value beyond AI. The results can inform FinTech companies and banks about staffing strategies and performance evaluation in an era of accelerating technological integration.

We would like to conclude with a few caveats. First, because loan officers make the final decision on loan applications, we do not observe the outcomes that would have happened had the company followed AI recommendations. Thus, we can only estimate the profit based on AI recommendations, and such estimation naturally contains errors. While we do not believe that the errors will introduce bias to our analyses, readers should keep this limitation in mind. Second, we focus on loans to small businesses. While the quantitative magnitude of our results might not be generalizable to the setting of large loans to corporations, where the decisions are more complicated and take a longer time to make, we believe that our qualitative findings on how the value of human intervention varies with the soft information used by loan officers and agency issues and behavioral bias faced by loan officers should apply to other settings. Lastly, our results might not be generalizable to future periods. AI technology is advancing fast, and the value of human intervention is likely to change accordingly. We leave these issues for future research.

References

Abis, S., and L. Veldkamp. 2024. The changing economics of knowledge production. *Review of Financial Studies* 37: 89–118.

Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo. 2022. Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics* 40 (S1): S293–S340.

Agrawal, A., J. S. Gans, and A. Goldfarb. 2019. Exploring the impact of artificial intelligence: Prediction versus judgment. *Information Economics and Policy* 47: 1–6.

Babina, T., A. Fedyk, A. He, and J. Hodson. 2024. Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics* 151: 103745.

Baumeister, R. E., E. Bratslavsky, M. Muraven, and D. M. Tice. 1998. Ego depletion: Is the active self a limited resource? *Journal of Personality and Social Psychology* 74 (5): 1252–1265.

Bertomeu, J., Y. Lin, Y. Liu, and Z. Ni. 2025. The impact of generative AI on information processing: Evidence from the ban of ChatGPT in Italy. *Journal of Accounting and Economics* 80 (1): 101782.

Bockstedt, J. C., and J. R. Buckman. 2025. Humans' use of AI assistance: The effect of loss aversion on willingness to delegate decisions. *Management Science* forthcoming.

Brynjolfsson, E., T. Mitchell, and D. Rock. 2018. What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings* 108: 43–47.

Bu, D., and Y. Liao. 2022. Shaming microloan delinquents: Evidence from a field experiment in China. *Management Science* 68 (5): 3768–3790.

Campbell, D., M. Loumioti, and R. Wittenberg-Moerman. 2019. Making sense of soft information: Interpretation bias and loan quality. *Journal of Accounting and Economics* 68 (2-3): 101240.

Cao, S., W. Jiang, J. Wang, and B. Yang. 2024. From man vs. machine to man + machine: The art and AI of stock analysts. *Journal of Financial Economics* 160: 103910.

Chu, J., R. Fisman, S. Tan, and Y. Wang. 2021. Hometown Ties and the Quality of Government Monitoring: Evidence from Rotation of Chinese Auditors. *American Economic Journal: Applied Economics* 13 (3): 176–201.

Coleman, B., K. Merkley, and J. Pacelli. 2022. Human versus machine: A comparison of robo-analyst and traditional research analyst investment recommendations. *The Accounting Review* 97 (5): 221–244.

Commerford, B. P., S. A. Dennis, J. R. Joe, and J. W. Ulla. 2022. Man versus machine: Complex estimates and auditor reliance on artificial intelligence. *Journal of Accounting Research* 60 (1): 171–201.

Commerford, B. P., A. Eilifsen, R. C. Hatfield, K. M. Holmstrom, and F. Kinserdal. 2024. Control issues: How providing input affects auditors' reliance on artificial intelligence. *Contemporary Accounting Research* 41 (4): 2134–2162.

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

Deng, Y., Z. Zhang, and Y. Liu. 2023. Auditors' hometown ties and audit quality. *Journal of Accounting and Public Policy* 42 (6): 107137.

Dietvorst, B. J., J. P. Simmons, and C. Massey. 2015. Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General* 144 (1): 114–126.

Fedyk, A., T. Fedyk, J. Hodson, and N.V. Khimich. 2022. Is artificial intelligence making audit firms more efficient? *Review of Accounting Studies* 27: 938–985.

Fisman, R., J. Shi, Y. Wang, and W. Wu. 2020. Social ties and the selection of China's political elite. *American Economic Review* 110 (6): 1752–1781.

Fuster, A., M. Plosser, P. Schnabl, and J. Vickery. 2019. The role of technology in mortgage lending. *Review of Financial Studies* 32 (5): 1854–1899.

Gerken, W. C., and M. O. Painter. 2023. The value of differing points of view: Evidence from financial analysts' geographic diversity. *Review of Financial Studies* 36 (2): 409–449.

Gopal, M., and P. Schnabl. 2022. The rise of finance companies and FinTech lenders in small business lending. *Review of Financial Studies* 35 (11): 4859–4901.

Grennan, J., and R. Michaely. 2021. FinTechs and the market for financial analysis. *Journal of Financial and Quantitative Analysis* 56 (6): 1877–1907.

Hirshleifer, D., Y. Levi, B. Lourie, and S. H. Teoh. 2019. Decision fatigue and heuristic analyst forecasts. *Journal of Financial Economics* 133 (1): 83–98.

Huettner, F., and M. Sunder. 2012. Axiomatic arguments for decomposing goodness of fit according to Shapley and Owen values. *Electronic Journal of Statistics* 6: 1239–1250.

Jansen, M., H. Q. Nguyen, and A. Shams. 2025. Rise of the machines: The impact of automated underwriting. *Management Science* 71 (2): 955–975.

Jussupow, E., I. Benbasat, and A. Heinzl. 2024. An integrative perspective on algorithm aversion and appreciation in decision-making. *MIS Quarterly* 48 (4): 1575–1590.

Kleinberg, J., H. Lakkaraju, J. Leskovec, J. Ludwig, and S. Mullainathan. 2018. Human decisions and machine predictions. *Quarterly Journal of Economics* 133 (1): 237–293.

Liberti, J., and A. Mian. 2009. Estimating the effect of hierarchies on information use. *Review of Financial Studies* 22: 4057–4090.

Liberti, J. M., and M. A. Petersen. 2019. Information: Hard and soft. *Review of Corporate Finance Studies* 8 (1): 1–41.

Liu, M. 2022. Assessing human information processing in lending decisions: A machine learning approach. *Journal of Accounting Research* 60 (2): 607-651.

McKinsey Analytics. 2025. The state of AI in 2025: Agents, innovation, and transformation. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai#/.

Petersen, M. A., and G. Rajan. 1994. The Benefits of lending relationships: Evidence from small business data. *The Journal of Finance* 49: 3–37.

Petersen, M. A., and G. Rajan. 2002. Does distance still matter? The information revolution in small business lending. *The Journal of Finance* 57: 2533–2570.

Stanovich, K. E., and R. F. West. 2000. Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences* 23 (5): 645–665.

Stein, J. 2002. Information production and capital allocation: Decentralized vs. hierarchical firms. *The Journal of Finance* 57 (5): 1891–1921.

Sutherland, A. G. 2020. Technology is changing lending: Implications for research. *Journal of Accounting Economics* 70 (2–3): 101361.

van Binsbergen, J. H., X. Han, and A. Lopez-Lira. 2023. Man versus machine learning: The term structure of earnings expectations and conditional biases. *Review of Financial Studies* 36 (6): 2361–2396.

APPENDIX A
Calculation of profit based on loan officers' decisions and AI's recommendations

In this appendix, we use an example to illustrate the calculation of profit of each loan based on the actual loan outcomes and AI's recommendations.

Panel A: Information about the loan and invoice

AI's recommendation: approval with a loan amount up to ¥100,000 at an annual interest rate of 15.8%

Loan officer's decision: loan amount up to ¥200,000 at an annual interest rate of 13.8%

Note that while both AI and loan officers make decisions at the loan application level, borrowing and payments occur at the invoice level.

Loan invoice ID: 000LIA2024010000067779

Invoice amount (the amount borrowed): ¥100,000

Borrowing date and time: 2024-01-01 14:04:35

Payment schedule for the invoice:

Payment Date and Time	Repaid Principal	Paid Interest
2024/2/1 8:17	3,678.16	1,066.60
2024/3/1 6:25	3,717.39	1,027.37
2024/4/1 8:00	3,757.04	987.72
2024/5/1 5:41	3,797.11	947.65
2024/6/1 12:28	3,837.61	907.15
2024/6/30 20:00	48,000.00	824.31
2024/7/1 8:22	1,928.79	11.62
2024/7/3 15:27	10,000.00	21.90
2024/8/1 8:23	1,089.80	216.05
2024/9/1 8:37	1,090.46	215.39
2024/10/1 9:22	1,102.09	203.76
2024/11/1 8:40	1,113.85	192.00
2024/12/1 8:28	1,125.73	180.12
2024/12/17 8:41	15,761.97	88.27

Note that interest payments are lower than the accrued interest based on the borrowed amount and the interest rate because of interest rate discounts.

Panel B: Calculation of Profit based on actual cash flows

For each payment, we calculate its present value by using the monthly average Shanghai Interbank Offered Rate (Shibor) for the corresponding period. *Profit* for the above invoice based on actual cash flows is calculated as follows:

$$\text{Profit}_{-\text{Man}} = \frac{3{,}678.16 + 1{,}066.6}{(1+\text{Shibor}_{24\text{-}1})^{31/365}} + \frac{3{,}717.39 + 1{,}027.37}{(1+\text{Shibor}_{24\text{-}1})^{31/365}(1+\text{Shibor}_{24\text{-}2})^{29/365}} + \ldots\ldots - 100{,}000 = 5{,}933.38$$

Shibor$_{24\text{-}1}$ refers to the average Shibor rate for January of 2024. Other Shibor variables are defined similarly.

Panel C: Calculation of Profit based on AI's recommendations

To calculate the hypothetical invoice amount and payments based on AI's recommendations, we calculate the adjustment ratio (*Adj_Ratio*) based on AI's recommended amount and loan officer's recommended amount as:

$$\text{Adj\_Ratio} = \text{AI Recommended Amount/Loan Amount by Loan Officer}$$

For the above example: $\text{Adj\_Ratio} = 100{,}000/200{,}000 = 1/2$

Because many loans have interest rate discount, we calculate the interest discount (*Discount*) for each loan in each month as:

$$\text{Discount}_m = \frac{\text{Paid Interest}_m/\text{Unrepaid\_Amount}_m}{(1+\text{Annual Interest Rate})^{\text{Time}/365}\text{-}1},$$

where *Time* denotes the number of days between the payment timestamp and the previous payment timestamp.

For the above example:

$$\text{Discount}_{24\text{-}1} = \frac{1066.6/100000}{(1+13.8\%)^{31/365}\text{-}1} = 96.61\%$$

$$\text{Discount}_{24\text{-}2} = \frac{1027.37/(100000\text{-}3678.16)}{(1+13.8\%)^{29/365}\text{-}1} = 99.85\%$$

...

We then use the *Adj_Ratio* and *Discount* to calculate the hypothetical AI-based borrowed amount, principal payment, and interest payment. Specifically, we calculate the hypothetical AI-based borrowed amount as:

$$\text{AI Invoice Amount} = \text{AI Recommended Amount} \times \textit{Adj\_Ratio},$$

the hypothetical AI-based repaid principal for each month as:

$$\text{AI Repaid Principal}_m = \text{Repaid Principal}_m \times \textit{Adj\_Ratio},$$

and the hypothetical AI-based interest payment for each month as:

$$\text{AI Paid Interest}_m = \text{Unpaid Amount}_m \times \textit{Adj\_Ratio}_m \times [(1+\text{AI Interest Rate})^{\text{Time}/365}\text{-}1] \times \textit{Discount}_m$$

Note that for the loans approved by AI, *AI Interest Rate* is the interest rate recommended by AI, and for the loans rejected by AI, *AI Interest Rate* is the average of the actual interest rate of the loans approved in the previous month.

For the above example:

$$\text{AI Invoice Amount} = 100{,}000 \times \frac{1}{2} = 50{,}000$$

$\text{AI Paid Principal and Interest}_{24\text{-}1} =$

$$3{,}678.16 \times \frac{1}{2} + 100{,}000 \times \frac{1}{2} \times [(1+15.8\%)^{\frac{31}{365}}\text{-}1] \times 96.61\% = 2{,}464.99$$

$\text{AI Paid Principal and Interest}_{24\text{-}2} =$

$$3{,}717.39 \times \frac{1}{2} + (100{,}000\text{-}3{,}678.16) \times \frac{1}{2} \times [(1+15.8\%)^{\frac{29}{365}}\text{-}1] \times 99.85\% = 2{,}461.58$$

...

We then calculate *Profit* using these values to obtain AI-based profit (*Profit*$_{\_AI}$). For the above example, $Profit_{\_AI}$ is calculated as:

$$\text{Profit}_{\_AI} = \frac{2{,}464{,}99}{(1+\text{Shibor}_{24\text{-}1})^{31/365}} + \frac{2{,}461.58}{(1+\text{Shibor}_{24\text{-}1})^{31/365}(1+\text{Shibor}_{24\text{-}2})^{29/365}} + \ldots - 50{,}000 = 3{,}491.05$$

Panel D: Calculation of the deviation in profit

For loans with multiple invoices, we aggregate $Profit_{\_Man}$ and $Profit_{\_AI}$ at the invoice level to obtain the profits at the loan level. We then calculate the deviation in profit of loan officers' decisions from AI recommendations as the difference in profit at the loan level ($Profit_{\_Man} - Profit_{\_AI}$), scaled by the total invoice amount for the loan.

Since the above example has one invoice, the deviation in profit is calculated as:

$$Deviation\_Profit = (5{,}933.38 - 3{,}491.05)/100{,}000 = 2.44\%, \text{ or}$$

$$Deviation\_Profit(\%) = 2.44$$

APPENDIX B
Variable Definitions

Variable	Definition
Loan profit variables	
Profit $_{\_Man}$	Actual profit of the loan based on the invoice-level cash flows; please see Appendix A for details of the calculation. For loans with multiple invoices, we aggregate the profit across invoices to obtain the loan-level profit.
Profit $_{\_AI}$	Hypothetical profit of the loan based on AI's recommendation; please see Appendix A for details of the calculation. For loans approved by AI, we use AI recommended interest rate; for loans rejected by AI, we use the average interest rate of the loans approved in the previous month. For loans with multiple invoices, we aggregate the profit across invoices to obtain the loan-level profit.
Deviation_Profit (%)	The difference in profit between loan officers' decision and AI's recommendation, calculated as $100 \times \frac{Profit_{-Man} - Profit_{-AI}}{Invoice\ Amount_{-Man}}$, where *Invoice Amount* $_{Man}$ is the total invoice amount, i.e., the total dispensed amount to the borrower (not the approved credit limit).
Loan outcome variables	
Loan Amount $_{\_Man}$	The final loan amount (i.e., credit limit) approved by the loan officer.
Interest Rate $_{\_Man}$	The final loan interest rate approved by the loan officer.
Loan Term $_{\_Man}$	The final loan term (in months) approved by the loan officer.
Invoice Amount $_{\_Man}$	The total invoice amount borrowed by the borrower under the loan.
Deviation_Amount	The difference in loan amount between the loan officer's decision and AI's recommendation, scaled by loan amount recommended by AI.
Deviation_Rate (%)	The difference in interest rate between the loan officer's decision and AI's recommendation times 100. For loans approved by AI, we use AI recommended interest rate; for loans rejected by AI, we use the average interest rate of the loans approved in the previous month.
Deviation_Term	The difference in term (in months) between the loan officer's decision and AI's recommendation.
Override	An indicator variable that equals one if the loan application is rejected by AI but approved by the loan officer, and zero otherwise.
Proxies for soft information	
Comment_Long	An indicator variable that equals one if the length of the loan officer's comment in the justification for loan decision is equal to or above the 75^{th} percentile of the sample distribution, and zero otherwise.
Business_Prospect	An indicator variable that equals one if the loan officer's comment in the justification for loan decision mentions the borrower's business prospects, and zero otherwise.
Family_Support	An indicator variable that equals one if the loan officer's comment in the justification for loan decision mentions the support for the borrower from his/her family members, and zero otherwise.
Credit_Anomalies	An indicator variable that equals one if the loan officer's comment in the justification for loan decision provides explanations for prior or current lawsuits or cash flow anomalies that might affect the borrower's creditworthiness, and zero otherwise.
Referral	An indicator variable that equals one if the borrower is recommended by existing clients or acquaintances, and zero otherwise.
Proxies for agency issues	

Hometown_Tie	An indicator variable that equals one if the borrower is from the same province as the loan officer, and zero otherwise.
Remote_Area	An indicator variable that equals one if the customer officer is located in a remote region as defined by the company (i.e., Yunnan, Guizhou, Gansu, or Shaanxi provinces) (customer officers in these regions are less likely to be visited by loan officers for periodic assessments), and zero otherwise.
Peer_Pressure	An indicator variable that equals one if, at the time when the loan is approved, the number of loans approved by the loan officer from the beginning of the day is lower than the median number of loans approved by all loan officers for the same period of the previous week, and zero otherwise.
Proxies for behavioral bias	
Multi_Loan	An indicator variable that equals one if the borrower has outstanding loans from other financial institutions at the time of loan application, and zero otherwise.
Infor_Request	An indicator variable that equals one if the loan officer requests additional documents (e.g., business licenses or bank statements) that have already been processed by AI, and zero otherwise.
Peak_Hour	An indicator variable that equals one if the loan is approved during daily peak approval hours (i.e., 10-11am and 4-5pm), and zero otherwise.
Around_Holiday	An indicator variable that equals one if the loan is approved around public holidays (i.e., the afternoon before or the morning after a public holiday, or on a compensatory working day), and zero otherwise.[24]

[24] The central government in China tends to arrange longer holidays but requires employees to work on the preceding and/or following weekends (referred to as compensatory working days). For example, in 2024, the government stipulated an eight-day holiday around Chinese New Year but required workers to work on the Sundays both right before and right after the holiday. The information on statutory holidays in 2024 and their compensatory working day arrangements can be downloaded from https://www.gov.cn/zhengce/content/202310/content_6911527.htm.

APPENDIX C
Correlation Matrix

This table reports the Pearson correlation matrix for the proxies for soft information, agency issues, and behavioral bias. All variables are defined in Appendix B. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Comment_Long (1)	1										
Business_Prospect (2)	0.144***	1									
Family_Support (3)	0.169***	0.036***	1								
Credit_Anomalies (4)	0.198***	0.023***	0.049***	1							
Referral (5)	0.051***	0.008	0.018**	0.010	1						
Hometown_Tie (6)	0.057***	0.031***	0.037***	0.027***	0.031***	1					
Remote_Area (7)	-0.022***	0.015**	0.002	-0.011	-0.064***	-0.087***	1				
Peer_Pressure (8)	0.002	0.014*	-0.020***	0.001	0.003	-0.002	0.010	1			
Multi_Loan (9)	0.004	0.003	-0.006	0.019***	-0.010	-0.001	0.018**	0.022***	1		
Infor_Request(10)	-0.039***	-0.188***	-0.177***	-0.194***	-0.335***	-0.020***	0.035***	0.007	0.036***	1	
Peak_Hour (11)	-0.021***	-0.008	-0.013*	-0.013*	0.006	-0.005	-0.002	0.005	-0.009	-0.025***	1
Around_Holiday (12)	-0.014*	0.001	0.004	-0.018**	0.008	-0.006	0.007	0.010	0.005	-0.006	0.002

FIGURE 1
Distribution of the Deviation in Profit

This figure depicts the distribution of the deviation in profit (*Deviation_Profit*). The x-axis indicates the value of *Deviation_Profit*, with intervals of 1%. The y-axis indicates the number of loans in each interval. The red dashed line indicates the zero interval, i.e., *Deviation_Profit* (%) in [0, 1).

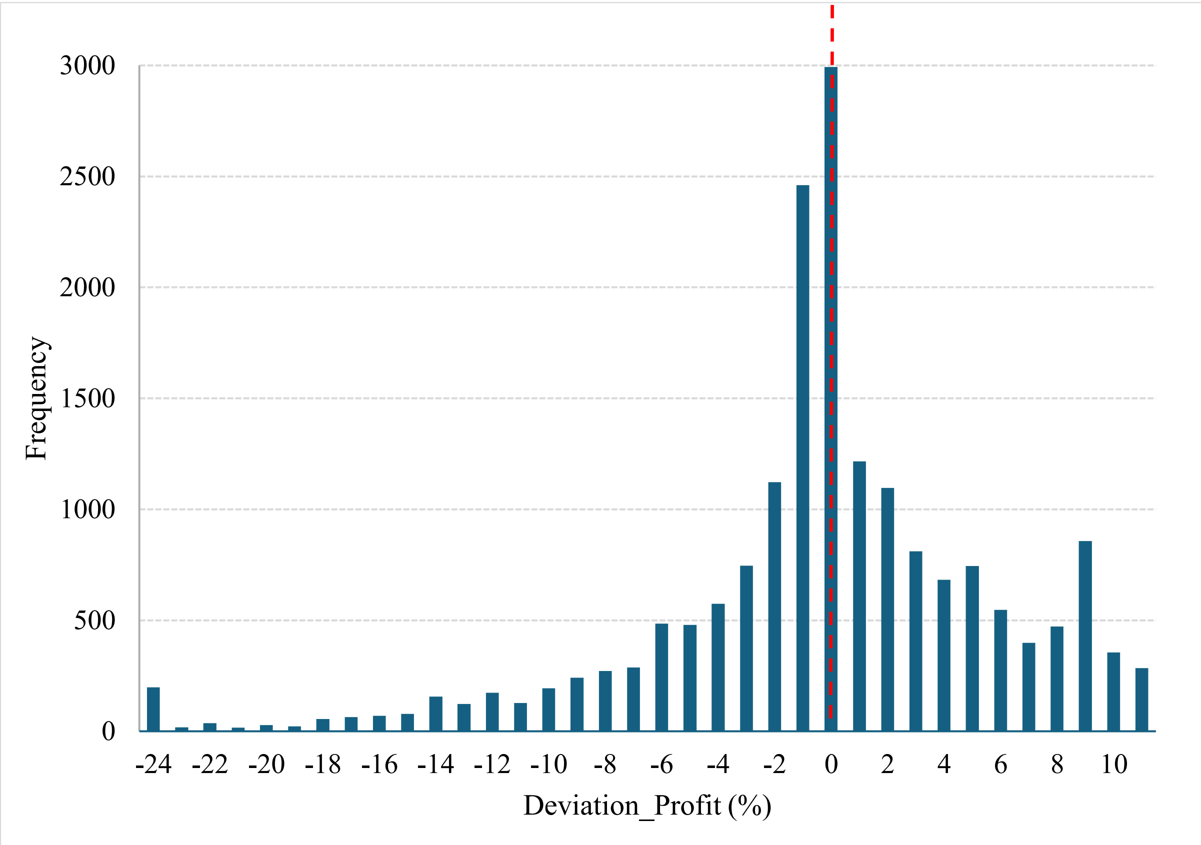

TABLE 1
Sample Selection

This table reports the sample selection process for the main analysis. The final sample consists of 18,490 loan observations from 2024.

	Obs.
All loans approved in 2024	28,198
Less:	
Loans that have invoices not yet matured at the time of data collection	(6,791)
Loans with missing values on loan outcomes or other variables	(2,706)
Loans approved by loan officers who approved fewer than 100 loans in 2024	(52)
Singleton observations	(159)
Final sample	18,490

TABLE 2
Descriptive Statistics on Loan Outcomes and Proxies for Soft Information, Agency Issues, and Behavioral Bias

This table reports descriptive statistics on loan outcomes and proxies for soft information, agency issues, and behavioral bias. Panel A reports the means and standard deviations of loan outcomes (loan amount, interest rate, and loan term) based on AI recommendations and loan officers' decisions, and the difference in means between AI recommendations and loan officers' decisions. Hypothetic AI Rate refers to the interest rate used in calculating loan profit under AI recommendation; it is set as the AI recommended interest rate for loans approved by AI and the average interest rate of the loans approved in the previous month for loans rejected by AI. Panel B reports descriptive statistics on the proxies for soft information, agency issues, and behavioral bias. Appendix B provides variable definitions. The sample consists of 18,490 loan observations from 2024.

Panel A: Loan outcomes based on AI recommendations and loan officers' decisions

	AI recommendations		Loan officers' decisions		Difference in mean	
	Mean (1)	Std. (2)	Mean (3)	Std. (4)	(3) − (1)	*t*-statistics
Loan Amount (CNY)	126,947	97,540	161,846	50,632	34,899***	46.21
Interest Rate (%)	18.15	1.60	14.51	2.54	-3.64***	-170.00
Hypothetical AI Rate (%)	16.39	1.85			-1.87***	-91.74
Loan Term (month)	16.58	6.28	24.62	11.37	8.04***	83.99

Panel B: Descriptive statistics on proxies for soft information, agency issues, and behavioral bias

	Mean	Std.	Q1	Median	Q3
Comment_Long	0.253	0.435	0	0	1
Business_Prospect	0.141	0.348	0	0	0
Family_Support	0.127	0.333	0	0	0
Credit_Anomalies	0.150	0.357	0	0	0
Referral	0.343	0.475	0	0	1
Hometown_Tie	0.050	0.217	0	0	0
Remote_Area	0.126	0.332	0	0	0
Peer_Pressure	0.411	0.492	0	0	1
Multi_Loan	0.383	0.486	0	0	1
Infor_Request	0.176	0.381	0	0	0
Peak_Hour	0.256	0.436	0	0	1
Around_Holiday	0.050	0.219	0	0	0

TABLE 3
Deviation in Profit: Descriptive Statistics and Determinants

This table reports descriptive statistics on and determinants of the deviation in loan profit. Panel A reports the means and standard deviations of *Profit* based on AI recommendations and loan officers' decisions, and the difference in mean between the two. Panel B reports descriptive statistics for the deviation in profit, loan amount, interest rate, and term, and the likelihood of loan officers overriding AI's recommendations of loan denial. Panel C reports the regression results of the deviation in profit (*Deviation_Profit (%)*) on the difference in loan outcomes—amount (*Deviation_Amount*), interest rate (*Deviation_Rate*), loan term (*Deviation_Term*)—and the likelihood of loan officers overriding AI's recommendations (*Override*). Intercepts are included but not tabulated. Appendix B provides variable definitions. The sample consists of 18,490 loan observations from 2024. For Panel C, the *t*-statistics in parentheses are based on standard errors adjusted for clustering at the loan officer level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on two-sided tests.

Panel A: Profit based on AI recommendations and loan officers' decisions

Profit (CNY)	AI recommendations (1)	Loan officers' decisions (2)	Difference in Mean (2) – (1)	*t*-statistic
Mean	12,109.10	13,153.45	1,044.36***	13.38
Std.	12,083.90	8,297.73		

Panel B: Descriptive statistics on the deviation in profit and loan outcomes

Variable	Mean	Std.	Q1	Median	Q3
Deviation_Profit (%)	0.209	6.164	-1.813	0.177	3.644
Deviation_Amount	2.134	3.478	-0.100	0.333	4
Deviation_Rate (%)	-1.875	2.779	-3.100	-1.700	-0.105
Deviation_Term	8.043	13.022	0	12	24
Override	0.609	0.488	0	1	1

Panel C: Determinants of the deviation in profit

Dependent variable =	*Deviation_Profit (%)*				
	(1)	(2)	(3)	(4)	(5)
Deviation_Amount	1.112***				1.132***
	(63.13)				(62.82)
Deviation_Rate (%)		0.393***			0.477***
		(20.14)			(29.08)
Deviation_Term			0.028***		0.020***
			(5.82)		(5.28)
Override				1.750***	-0.010
				(14.60)	(-0.11)
Month FE	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes
Loan Officer FE	Yes	Yes	Yes	Yes	Yes
Obs.	18,490	18,490	18,490	18,490	18,490
Adj. R^2	0.374	0.039	0.014	0.027	0.418

TABLE 4
Soft Information and Deviation in Profit

This table reports the regression results of the deviation in profit (*Deviation_Profit*) on the proxies for soft information used by loan officers. Appendix B provides variable definitions. Intercepts are included but not tabulated. The sample includes 18,490 loan observations in 2024. The *t*-statistics in parentheses are based on standard errors adjusted for clustering at the loan officer level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on two-sided tests.

Dependent variable =	*Deviation_Profit (%)*				
	(1)	(2)	(3)	(4)	(5)
Comment_Long	0.366***				
	(2.65)				
Business_Prospect		0.215**			
		(2.05)			
Family_Support			0.460***		
			(2.92)		
Credit_Anomalies				0.440***	
				(2.95)	
Referral					0.219***
					(2.70)
Month FE	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes
Loan Officer FE	Yes	Yes	Yes	Yes	Yes
Obs.	18,490	18,490	18,490	18,490	18,490
Adj. R^2	0.011	0.011	0.011	0.011	0.011

TABLE 5
Agency Issues and Deviation in Profit

This table reports the regression results of the deviation in profit (*Deviation_Profit*) on the proxies for agency issues faced by loan officers. Appendix B provides variable definitions. Intercepts are included but not tabulated. The sample includes 18,490 loan observations in 2024. The *t*-statistics in parentheses are based on standard errors adjusted for clustering at the loan officer level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on two-sided tests.

Dependent variable =	*Deviation_Profit (%)*		
	(1)	(2)	(3)
Hometown_Tie	-0.998***		
	(-4.44)		
Remote_Area		-0.596***	
		(-4.70)	
Peer_Pressure			-0.700***
			(-5.98)
Month FE	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes
Loan Officer FE	Yes	Yes	Yes
Obs.	18,490	18,490	18,490
Adj. R^2	0.012	0.014	0.014

TABLE 6
Behavioral Bias and Deviation in Profit

This table reports the regression results of the deviation in profit (*Deviation_Profit*) on the proxies for behavioral bias of loan officers. Appendix B provides variable definitions. Intercepts are included but not tabulated. The sample includes 18,490 loan observations in 2024. The *t*-statistics in parentheses are based on standard errors adjusted for clustering at the loan officer level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on two-sided tests.

Dependent variable =	*Deviation_Profit (%)*			
	(1)	(2)	(3)	(4)
Multi_Loan	-1.465***			
	(-12.07)			
Infor_Request		-0.481***		
		(-3.52)		
Peak_Hour			-0.233***	
			(-2.68)	
Around_Holiday				-0.385***
				(-2.60)
Month FE	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes
Loan Officer FE	Yes	Yes	Yes	Yes
Obs.	18,490	18,490	18,490	18,490
Adj. R^2	0.024	0.011	0.011	0.011

TABLE 7
Analyses of the Deviation in Profit Separately for Loans Approved in the First and Second Half of 2024

This table reports the regression results of the deviation in profit (*Deviation_Profit*) on the proxies for soft information used by loan officers (Panel A), agency issues faced by loan officers (Panel B), and behavioral bias of loan officers (Panel C), separately for loans approved in the first half of 2024 and in the second half of the year. We use the specifications in Tables 4-6 to estimate the effect of each proxy on *Deviation_Profit* (%). For brevity, the table only reports the coefficient and *t*-statistics for the variables of interest. The bottom of each panel reports the differences in coefficients between the two subsamples and the corresponding *p*-values in parentheses. Appendix B provides variable definitions. Intercepts are included but not tabulated. The sample includes 18,490 loan observations in 2024, with 11,375 observations in the first half and 7,115 in the second half. The *t*-statistics in parentheses are based on standard errors adjusted for clustering at the loan officer level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on two-sided tests.

Panel A: Soft information and the deviation in profit

Dependent variable =	*Deviation_Profit (%)*				
Soft information proxy =	*Comment_Long*	*Business_Prospect*	*Family_Support*	*Credit_Anomalies*	*Referral*
Loans approved in January-June of 2024					
Soft information proxy (1)	0.263	0.022	0.204	0.404**	0.267***
	(1.55)	(0.16)	(0.87)	(2.25)	(3.06)
Loans approved in July-December of 2024					
Soft information proxy (2)	0.479***	0.466***	0.787***	0.483***	0.154
	(2.97)	(2.84)	(4.18)	(3.01)	(1.16)
Difference in coefficient: (2) – (1)	0.216	0.444*	0.583**	0.079	-0.113
(p-value)	(0.184)	(0.062)	(0.020)	(0.438)	(0.300)

TABLE 7 (cont'd)

Panel B: Agency issues and the deviation in profit

Dependent variable =	*Deviation_Profit (%)*		
Agency issue proxy =	*Hometown_Tie*	*Remote_Area*	*Peer_Pressure*
Loans approved in January-June of 2024			
Agency issue proxy (1)	-1.569***	-0.884***	-0.597***
	(-6.46)	(-5.00)	(-4.63)
Loans approved in July-December of 2024			
Agency issue proxy (2)	-0.306	-0.203	-0.937***
	(-0.87)	(-1.52)	(-4.25)
Difference in coefficient: (2) – (1)	1.263***	0.681***	-0.340**
(p-value)	(0.01)	(0.002)	(0.038)

Panel C: Behavioral bias and the deviation in profit

Dependent variable =	*Deviation_Profit (%)*			
Cognitive constraint proxy =	*Multi_Loan*	*Infor_Request*	*Peak_Hour*	*Around_Holiday*
Loans approved in January-June of 2024				
Cognitive constraint proxy (1)	-1.324***	-0.353***	-0.001	-0.221
	(-8.53)	(-2.82)	(-0.01)	(-0.96)
Loans approved in July-December of 2024				
Cognitive constraint proxy (2)	-1.718***	-0.737***	-0.603***	-0.846***
	(-12.04)	(-2.91)	(-4.58)	(-2.61)
Difference in coefficient: (2) – (1)	-0.394**	-0.383*	-0.602***	-0.625
(p-value)	(0.02)	(0.076)	(0.004)	(0.108)

TABLE 8
The Relative Importance of Soft Information, Agency Issues, and Behavioral Bias

This table reports the regression results of the deviation in profit (*Deviation_Profit*) on the proxies for soft information used by loan officers, agency issues faced by loan officers, and behavioral bias of loan officers in one regression. Appendix B provides variable definitions. Intercepts are included but not tabulated. The sample includes 18,490 loan observations in 2024. The *t*-statistics in parentheses are based on standard errors adjusted for clustering at the loan officer level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on two-sided tests.

Dependent variable =		*Deviation_Profit (%)*			
		(1)	(2)	(3)	(4)
Soft information	*Comment_Long*	0.246*			0.294**
		(1.88)			(2.22)
	Business_Prospect	0.178*			0.178
		(1.66)			(1.56)
	Family_Support	0.405**			0.358**
		(2.49)			(2.14)
	Credit_Anomalies	0.372***			0.338**
		(2.63)			(2.06)
	Referral	0.215***			0.158*
		(2.61)			(1.75)
Agency issues	*Hometown_Tie*		-1.062***		-1.167***
			(-4.52)		(-4.83)
	Remote_Area		-0.635***		-0.564***
			(-4.83)		(-4.39)
	Peer_Pressure		-0.694***		-0.644***
			(-5.87)		(-5.82)
Behavioral bias	*Multi_Loan*			-1.453***	-1.443***
				(-12.16)	(-12.27)
	Infor_Request			-0.424***	-0.196
				(-3.12)	(-1.30)
	Peak_Hour			-0.255***	-0.243***
				(-2.74)	(-2.62)
	Around_Holiday			-0.367**	-0.341**
				(-2.37)	(-2.20)
Month FE		Yes	Yes	Yes	Yes
Industry FE		Yes	Yes	Yes	Yes
Loan Officer FE		Yes	Yes	Yes	Yes
Obs.		18,490	18,490	18,490	18,490
Adj. R^2		0.012	0.016	0.025	0.030

TABLE 9
Analyses of the Deviation in Profit Using Alternative Design Choices

This table reports the regression results of the deviation in profit (*Deviation_Profit*) on the proxies for soft information, agency issues , and behavior bias using alternative design choices. In Column (1), the hypothetical profit for the loans approved by AI remains the same, but for the loans denied by AI, we calculate the hypothetical profit using the interest rate recommended by AI. In Column (2), we use the same value of *Deviation_Profit* but restrict the sample to loans not in default. In Column (3), the hypothetical profit for the loans under AI recommendations is calculated using an alternative measure of invoice amount; we use the lower of the actual invoice amount and the credit limit approved by AI. Payments of interest and principal are adjusted accordingly. Appendix B provides other variable definitions. Intercepts are included but not tabulated. The sample includes 18,490 loan observations in 2024. The *t*-statistics in parentheses are based on standard errors adjusted for clustering at the loan officer level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on two-sided tests.

Dependent variable =	*Deviation_Profit (%)*		
	(1) Alternative interest rate for $Profit_{AI}$	(2) Sample of loans not in default	(3) Alternative invoice amount for $Profit_{AI}$
Comment_Long	0.503***	0.189*	0.167
	(3.01)	(1.78)	(1.26)
Business_Prospect	0.317**	0.182	0.097
	(2.38)	(1.31)	(1.18)
Family_Support	0.338*	0.327*	0.133
	(1.82)	(1.80)	(1.25)
Credit_Anomalies	0.487***	0.334*	0.184**
	(2.90)	(1.93)	(1.98)
Referral	0.116	0.091	0.144*
	(0.86)	(0.89)	(1.82)
Hometown_Tie	-1.233***	-1.085***	-0.890***
	(-4.39)	(-4.37)	(-6.24)
Remote_Area	-0.713***	-0.640***	-0.411***
	(-3.72)	(-4.91)	(-4.65)
Peer_Pressure	-0.750***	-0.684***	-0.408***
	(-5.41)	(-6.76)	(-5.44)
Multi_Loan	-1.502***	-1.459***	-0.570***
	(-8.94)	(-11.00)	(-9.73)
Infor_Request	-0.167	-0.209	-0.123
	(-0.90)	(-1.18)	(-1.30)
Peak_Hour	-0.266**	-0.224**	-0.099
	(-2.50)	(-2.29)	(-1.59)
Around_Holiday	-0.300	-0.330*	-0.021
	(-1.34)	(-1.76)	(-0.18)
Month FE	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes
Loan Officer FE	Yes	Yes	Yes
Obs.	18,490	15,513	18,490
Adj. R^2	0.024	0.028	0.026

TABLE 10
Analyses of the Actual Profit

This table reports the regression results of the actual profit (*Profit*$_{Man}$) on AI profit (*Profit*$_{AI}$) and the proxies for soft information, agency issues, and behavior bias. Appendix B provides variable definitions. Intercepts are included but not tabulated. The sample includes 18,490 loan observations in 2024. The *t*-statistics in parentheses are based on standard errors adjusted for clustering at the loan officer level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on two-sided tests.

Dependent variable =		*Ln(Profit*$_{Man}$*)*
Ln(Profit$_{AI}$*)*		0.663***
		(64.67)
Soft information	*Comment_Long*	0.126***
		(5.36)
	Business_Prospect	0.041*
		(1.79)
	Family_Support	0.029
		(1.41)
	Credit_Anomalies	0.079***
		(3.45)
	Referral	0.067***
		(4.07)
Agency issues	*Hometown_Tie*	-0.000
		(-0.00)
	Remote_Area	-0.138***
		(-7.14)
	Peer_Pressure	-0.092***
		(-5.40)
Behavior bias	*Multi_Loan*	-0.141***
		(-13.09)
	Infor_Request	-0.005
		(-0.26)
	Peak_Hour	-0.015
		(-1.20)
	Around_Holiday	-0.005
		(-0.15)
Month FE		Yes
Industry FE		Yes
Loan Officer FE		Yes
Obs.		18,490
Adj. R^2		0.657