

Artificial intelligence and platform credit risk^{*}

Long Chen

Luohan Academy

Jon Frost

BIS

Yi Huang

BIS, CEPR

Yingju Ma

University of Cambridge

August 10, 2026

Preliminary draft, please do not circulate without authors' permission

Abstract

Breakthroughs in artificial intelligence (AI) are reshaping credit markets by enabling real-time risk management, not merely better screening at origination. We study what AI does, and what it cannot do, in unsecured small-business lending. Using proprietary, high-frequency data on a sample of 800,000 small-business borrowers drawn from a large platform lender, covering January 2018 to March 2021, we show that the lender's machine-learning-based credit system continuously adjusts credit access, credit lines, and interest rates as risk conditions change, including during the Covid-19 shock, when delinquencies on these unsecured loans rose but remained contained. Prediction alone, however, does not enforce repayment. Employing an instrumental-variables design, we identify a causal effect of platform dependence on default: greater dependence lowers default, consistent with stronger repayment incentives from continued platform participation. These findings reveal a distinctive feature of AI-powered platform lending: it is most effective when algorithmic risk management is integrated with platform access and activity.

Keywords: artificial intelligence, credit risk, lending, digital platforms, big tech, fintech, big data, algorithms. **JEL classifications:** E51, G23, O31.

^{*}Long Chen is with Luohan Academy. Jon Frost is with Bank for International Settlements (BIS), email: jon.frost@bis.org. Yi Huang is with BIS and CEPR, email: yi.huang@bis.org. Yingju Ma is with the University of Cambridge, email: ym459@cam.ac.uk. We benefited from the comments of Bengt Holmstrom, Patrick Bolton, Wei Xiong, Nan Li, Robert M. Townsend, and participants at a BIS research seminar. We thank the data support teams of Ant Group and MYBank, which have helped create de-identified data that are selectively available to be used for academic research. Zhiren Hu, Boyang Zhang, and Giulio Cornelli provided outstanding research assistance in the early stages of the project. The views expressed in this paper are those of the authors and do not necessarily represent those of BIS or Ant Group. For comments and suggestions, please email ym459@cam.ac.uk. (First posted version: October 30, 2024.)

1 Introduction

Advances in artificial intelligence (AI)¹ have accelerated dramatically in the past few years. The rapid diffusion of AI across finance, from credit scoring and asset management to debt collection and relationship lending, is reshaping how financial intermediaries originate, price, and monitor loans (Foucault et al., 2025; Huang, Lee, and Yeung, 2025; Chen, Sialm, and Xu, 2025). One particularly active set of lenders deploying AI are large digital platform firms (“big techs”). These digital platform providers are increasingly branching out into finance (Bank for International Settlements, 2019; Frost et al., 2019; Stulz, 2019). In some economies, AI-powered platform lenders have developed significant lending activities, either directly or in partnership with financial institutions. Such AI-driven platform credit is estimated to have reached over USD 500 billion globally in 2019 (Cornelli et al., 2023). The largest market for AI-powered platform credit by far is China, where big techs like Ant Group, Tencent, JD, Baidu and Meituan offer a growing range of financial services to individuals (eg, buyers on their platforms) and to small and medium-sized enterprises (SMEs).

As a new form of financial intermediation, such AI-powered digital platform credit may promote financial inclusion and inclusive growth (Luohan Academy, 2019). Digital pay-as-you-go financing models, enabled by AI and real-time data, have been shown to generate substantial welfare gains for credit-constrained households and small firms (Gertler et al., 2025), and the value of data itself for credit allocation and welfare is increasingly well documented (Jansen et al., 2025). Yet as with all forms of lending, it is subject to specific financial risks, namely credit risk (the risk that borrowers will default), liquidity risk (the risk that the lender will not be able to meet cash flow needs, eg to repay funding), operational risk (the risk of loss resulting from inadequate or failed internal pro-

¹Throughout the paper, we use the term AI in the sense of Agrawal, Gans, and Goldfarb (2018): machine-learning algorithms that turn large volumes of data into predictions, coupled with automated decisions based on those predictions. In lending, the prediction is a borrower’s default risk, and the decisions are whether to lend, how much, and at what price. This definition covers the machine-learning credit models studied in the recent literature (Fuster, Goldsmith-Pinkham, et al., 2022).

cesses, people and systems) and aggregate systematic risk. While these risks have been studied in detail for bank lending, there have been few analyses yet of financial risks for AI-based platform credit. In this project, we focus specifically on credit risk, assessing the impact of AI and specific features of digital platforms.

Aided by big data, emerging research suggests both advantages and risks of digital (fintech and big tech) lending (see [Berg, Fuster, and Puri \(2022\)](#)). Based on platform information advantages ([Agarwal et al., 2020](#)) and more efficient information acquisition from a digital payment system ([Ghosh, Vallee, and Zeng, 2025](#)), AI-powered platform credit can serve previously un(der)served segments of the population, with limited account access or credit history. It can also reduce geographic frictions in credit markets ([Gopal and Schnabl, 2022](#); [Huang, Zhang, et al., 2020](#); [Hau et al., 2019](#)). AI-based lending models can represent a technological innovation that reduces frictions in lending, such as lengthy loan processing, capacity constraints, inefficient refinancing and limited access to finance by some borrowers. These efficiency gains do not lead to higher default costs ([Fuster, Plosser, et al., 2019](#)). Moreover, AI systems can leverage big data as a substitute for physical collateral, providing inclusive finance to borrowers who lack traditional assets ([Gambacorta, Huang, et al., 2023](#)). Recent evidence further shows that AI is transforming relationship lending, enabling lenders to combine hard algorithmic signals with soft information in ways that improve credit allocation ([Gambacorta, Sabatini, and Schiaffi, 2025](#); [Chen, Cheng, et al., 2026](#)). Information production by lenders, including AI-augmented screening, also varies over the business cycle, with important implications for credit supply ([Howes and Weitzner, 2025](#)).

However, as with any innovation or new business model, the platform's information advantage could also lead to unintended consequences. Recent research shows that digital platforms could aggregate consumers' purchase histories into scores that proxy for unobserved willingness to pay. Platforms could take advantage of those hidden scores—which only they observe—to reduce demand sensitivity, increase average prices and re-

duce consumer surplus ([Bonatti and Cisternas, 2020](#)). Algorithmic decision-making could mitigate in-person discrimination bias, but it can also lead to new biases and discrimination, as shown by [Buchak et al. \(2018\)](#), [Blattner and Nelson \(2021\)](#), [Bhutta, Hizmo, and Ringo \(2022\)](#), [Di Maggio and Yao \(2021\)](#), [Fuster, Goldsmith-Pinkham, et al. \(2022\)](#) and [Bartlett et al. \(2022\)](#) show. Emerging evidence on AI-assisted debt collection further illustrates both the promise and the limits of algorithmic enforcement in credit markets ([Choi et al., 2025](#)).

Can AI systems, leveraging real-time big data, improve human decision-making to lend to borrowers within a platform? [Di Maggio, Ratnadiwakara, and Carmichael \(2022\)](#) show that the lenders can leverage big data to find “invisible primes”—borrowers with low credit scores and short credit histories, but also a low propensity to default. [Berg, Burg, Gombović, et al. \(2020\)](#), [Berg, Fuster, and Puri \(2022\)](#) and [Berg, Burg, Keil, et al. \(2022\)](#) exploit data from platform firms, especially payment firms, and study digital footprint variables (eg device type, the operating system and the email provider) to be important predictors of default. [Dai et al. \(2022\)](#) show that information acquired by the lender through borrowers’ digital footprints increases the repayment likelihood on delinquent loans. In addition, the lender tends to approve loan applications from borrowers with digital footprints, even though these borrowers occasionally have a higher likelihood of delinquency. Those findings suggest that digital footprints, as a new type of enforcement mechanism, can enhance financial inclusion.

More generally, our paper is related to recent studies focused on AI-powered lending models and platform ecosystems (see, for example, [Liu, Lu, and Xiong \(2022\)](#) and [Brunnermeier and Payne \(2022\)](#)). We contribute to this strand of the literature with evidence that AI-powered platform lending can be successfully employed not only to predict default but also for dynamic credit risk management. Those advantages for risk management rely on real-time (alternative) data and platform enforcement, given the threat of platform exclusion, a mechanism formalized by [Li and Pegoraro \(2026\)](#), who show that a

bigtech platform’s ability to exclude borrowers from its ecosystem strengthens repayment incentives. The growing deployment of AI across financial services—from credit origination to asset management ([Chen, Sialm, and Xu, 2025](#))—underscores the importance of understanding how these systems manage risk in practice.

This paper asks what AI does in lending beyond screening at origination, and what it cannot do. We study small and medium-sized enterprise (SME) lending by a large platform lender in China, whose credit decisions are made by a machine-learning-based system with no human intervention. Prior work has shown that access to platform-based finance can reduce firm-level volatility ([Chen, Huang, et al., 2022](#)). We use the Covid-19 shock, which affected different regions and borrowers differently, to shed light on how AI-based platform credit performs during shocks, and how its credit risk differs from that of traditional banks ([Vasicek, 1987](#)). We organise the evidence in three steps.

First, we document the credit-risk outcomes of the AI-based lending system. Delinquency rates are lower than those of banks that lend to comparable small borrowers. We confirm that the probability of default (PD) and loss given default (LGD) rose during the pandemic, particularly in cities more affected by the pandemic, or with tighter trade links to Wuhan. Yet notably, the increase was largely in the first months of the pandemic; by June 2020, the rise in delinquencies was relatively contained, and correlations started coming down again. This is particularly surprising given that such lending does not rely on collateral.

Second, we show how AI operates as a real-time risk management tool. The system adjusts credit admission, credit lines and interest rates as risk conditions change. It does so by reallocating credit across borrowers rather than by contracting uniformly, and it adjusts quantities far more frequently than prices. Because algorithmic decisions raise concerns about bias, we also examine whether these adjustments differ by borrowers’ gender, education and age.

Third, we show the limit of AI. Prediction addresses the lender’s information problem,

but it cannot by itself enforce repayment of an unsecured loan. Using an instrumental-variables design, we show that borrowers who depend more on the platform default less, and that this effect is concentrated among sellers with larger platform-specific investments. The value of continued platform participation thus acts as implicit collateral, given the threat of platform exclusion.

We go beyond the existing literature on AI-powered platform credit by identifying two channels through which the characteristics of such lending, as part of an integrated digital ecosystem and using real-time alternative (big) data, can mitigate credit risk. The first is dynamic credit adjustment by the AI system, which also helped offline borrowers. The second is the value of the digital platform services, which can entice borrowers, especially online borrowers, to continue repaying. Our evidence on the first channel is descriptive, while our evidence on the second is causal. To our knowledge, ours is the first study to show in a systematic fashion how specific features of AI-powered lenders affect the risk of platform credit portfolios, and how credit risk was managed in practice during the Covid shock. Finally, our research findings serve as an input to evidence-based policymaking on the regulation of AI-driven platform financial activities.

The rest of this paper is organised as follows. Section 2 describes the AI-based lending system. Section 3 documents its credit-risk outcomes: relative to banks, across cities, and during the Covid-19 shock. Section 4 studies AI as a real-time risk management tool, using aggregate and borrower-level evidence on dynamic credit adjustment, and examines algorithmic bias. Section 5 turns to the limits of AI and shows that platform dependence sustains the incentives to repay. Finally, Section 6 concludes.

2 The AI-based Lending System

We use the term AI to refer to the lender’s machine-learning-based credit system: a set of models that predict each borrower’s default risk from payment, transaction and be-


havioural data, combined with automated rules that map these predictions into decisions on credit admission, lines of credit and interest rates. Two features distinguish this system from conventional credit scoring. First, decisions are fully automated, with no human intervention. Second, predictions and decisions are not made once, at the time of a loan application; they are updated continuously, before and after origination. This section describes the lender, its borrowers and the architecture of the system. The following sections study what the system does and where its limits lie.

The firm is a major lender to small businesses across China. In this study, we focus on its two major lending products. One is provided to “online” e-commerce merchants, who sell products on a digital platform whose services are part of the same overall ecosystem of the platform lender, and use the lender’s payment functionality. The other lending product is for “offline” merchants—small business owners who use the firm’s digital payment functionality in their businesses, such as restaurants and grocery stores, but do not rely on the e-commerce platform. In total, there were 45 million borrowers as of February 2021, with 27 million loans, an average loan size of RMB 1,780 (~USD 280), an average duration of 1.5 months and an annualised interest rate ranging from 6% to 17%.

These activities grew fast in the period leading up to, and during, the Covid-19 pandemic. Strikingly, lending by the firm nearly doubled in 2020 relative to the previous year (Figure 1, left-hand panel). Overall platform lending in China also grew in 2020, by about 50% relative to 2019. Lending values thus reached about RMB 4400 bn (~USD 620 bn) (right-hand panel).

A distinctive feature of the platform’s lending model is that risk management begins at the *potential borrower* stage—well before any loan application—and continues through approval, drawdown, and repayment, with credit limits and pricing updated in real time as new data arrive. We now describe the risk management architecture that underlies the platform’s SME lending business, which enables continuous, data-driven adjustment of

Figure 1: Big tech credit kept growing even during the Covid-19 pandemic



Note: “Total big tech lending” refers both to lending directly by big techs and lending in partnership with financial institutions. Sources: authors’ calculations; [Cornelli et al. \(2023\)](#).

credit limits and pricing.²


2.1 The Risk Management Architecture

The system spans the full borrower lifecycle, starting from the screening of prospective SMEs in the platform’s merchant ecosystem. Transaction, behavioral, and network data are continuously ingested to pre-compute eligibility and indicative credit limits, so that once a borrower applies, approval is near-instantaneous. Figure 2 illustrates the architecture.

The credit risk management system consists of three sets of decisions. First is *credit admission*. With the vast customer base of the e-commerce platform and the digital payment service, the lender can evaluate the creditworthiness of a potential SME applicant and decide whether she will be granted a loan if she applies. This can be done for both online borrowers (with a wealth of e-commerce data) and offline borrowers (with only payment

²The platform’s SME lending product is often referred to as the “3-1-0 model”: less than three minutes for a borrower to apply, less than one second for the lender to approve, and zero human intervention. This fully automated process is made possible by a real-time, AI-based credit risk management system.

Figure 2: The credit risk management system behind AI-powered platform SME loans




data). If an SME owner is admitted, the second set of decisions on the size of the *line of credit* and *interest rates* are made by the pre-credit management system. To show a new applicant these decisions within one second, the pre-credit management system maintains the credit decisions for all potential borrowers and updates them at a high frequency (usually daily). Once a borrower applies for a loan, she will sign a borrowing contract with the lender, which allows her to apply for loans at any time as long as her loan request does not exceed the remaining line of credit. Finally, the credit adjustment system evaluates the borrower's creditworthiness and the history of repayments, and updates these at a regular frequency. The borrower will see the latest line of credit and interest rates on her end at any time.

What distinguishes this architecture from traditional credit systems is the frequency and granularity of adjustment. The system responds to aggregate risk conditions by tightening total credit supply when system-wide risk rises; it responds to borrower-level demand by raising temporary credit lines during peak sales seasons; and it responds to local shocks by scaling back credit in regions where the real economy is affected (Section 3.2).

Comparing online and offline borrowers sheds light on both the challenges of risk control and the AI system's capabilities. Online borrowers are SME owners who operate e-

commerce shops, food delivery services, or similar establishments on the platform; offline borrowers are small vendors who use only the platform's QR-code payment service. Loans to both groups have six- to twelve-month maturities with flexible prepayment—borrowers pay interest only over the period credit is drawn. Figure 3 shows that average loan size has grown over time to RMB 18,000 (~USD 2,800) for online borrowers and RMB 12,000 for (~USD 1,880) for offline borrowers in 2021, with medians below RMB 5,000 for both. On average, loans are repaid in one to two months.

Figure 3: Loans are typically small, but loan size has grown over time



Note: Authors' calculations.

Both groups depend on the platform, but to different degrees. Online borrowers rely on the associated e-commerce platform for sales as well as payments, and sellers who have accumulated customer reviews and reputation on the platform attach substantial value to remaining in good standing—even when financially strained. Offline borrowers, by contrast, use only the payment service and could in principle switch to a competing payment provider. Given the platform's significant share in the payment market, such a switch would cause inconvenience to a considerable number of consumers, which could in turn impact the borrowers' businesses. As we will discuss in more depth, the fact that offline borrowers continued to repay during the Covid shock could initially be seen as a

puzzle, and relates to specific practices aimed at dynamically adjusting such credit.

3 Credit Risk of AI-based Lending: Aggregate and Regional Evidence


In this section, we document the credit-risk outcomes of the AI-based lending system described in Section 2. We compare its lending performance with the overall risk level of the financial system and with city-level variations in credit risk across the economy, and we examine how the portfolio performed during the regional Covid-19 shock. The evidence in this section is descriptive: it shows how credit risk evolved under the AI-based system, and motivates the analysis of the mechanisms in Sections 4 and 5.

3.1 Delinquency Rates: Platform Lender versus Commercial Banks

We first calculate the probability of default (PD) for loans by the platform firm across its entire portfolio, at the national level. At any given point in time, the PD is defined as the ratio of the value of loans that are delinquent for 30 days or more to the total value of outstanding loans. We compare this with data major categories of Chinese banks.

Figure 4 shows the sampled AI-based platform loans' PD rate and that of different categories of banks in China. The PD of the platform loans is significantly lower than that of city commercial banks and rural commercial banks. These two types of banks issue loans with relatively smaller sizes than other banks, although still much larger than the platform loans. During the first half year of 2020, the risk of platform loans rises, reflecting the shock to SMEs due to the lockdown policies in China. It decreases after July 2020, consistent with the AI-powered lender quickly adjusting to contain credit risk. The default rates on joint stock commercial banks and the national average remain fairly constant over the full period.

Figure 4: Delinquency rates of fintech loans are lower than at commercial banks



Note: All the delinquency rates refer to a 90-day delinquency rate. Different types of banks are categorized based on the classification of commercial banks by the People’s Bank of China. Data is calculated by the authors.

3.2 Platform Credit Risk and Local Real-economy Defaults

To examine the relationship between AI-based credit risk and other forms of credit risk in the real economy, we introduce city-level data on defaults. In China, when a debt default occurs, if the debtor fails to fulfil the repayment obligation within a certain period following a court ruling, the court designates the individual as a defaulter and makes this information publicly available. Being listed entails various penalties, most notably restrictions on consumption activities. We collect data from the nationwide legal judgment system across all cities in the years 2018–20, and construct the annual number of defaulters for each city. In addition, using our daily city-level PD data, we compute the annual average PD for each city, based on both a 30-day definition and a stricter definition of 90 days past due. We also compile data on GDP per capita and GDP growth for the cities over this period.

Table 1 gives descriptive statistics. Notably, the mean of the city-level PD of the platform lender is about 1.92% with the 30-day definition, and 1.34% with the 90-day definition over the sample period. Meanwhile, from the nationwide legal judgment system (ie

defaults on loans from all sources), there are approximately four defaulters per thousand residents per year in the average city.

Table 1: Descriptive statistics at city level

Variable name	Mean	P25	Median	P75	Std. Dev.	N
PD (30 days)	1.92	1.23	1.83	2.54	0.89	1,046
PD (90 days)	1.34	0.76	1.26	1.83	0.72	1,046
Defaulters (per thousand population)	3.99	1.30	3.11	5.58	3.81	818
GDP per capita (RMB)	62,774	38,026	52,344	78,545	33,853	818
GDP growth rate (%)	5.33	3.70	5.90	7.60	3.27	817

Note: The table presents summary statistics for Fintech credit risk and real-economy risk, based on data from 349 cities during 2018–2020. The real-economy default data are available for 273 of these cities.

Table 2: AI-based credit risk at the city level

Dependent variable: probability of default, Fintech loans (30 days, percentage)					
	(I)	(II)	(III)	(IV)	(V)
Defaulters (per thousand)	0.024** (0.008)	−0.013* (0.005)	−0.023* (0.011)	−0.016** (0.006)	−0.027* (0.012)
GDP per capita				−0.152** (0.048)	−0.916*** (0.240)
GDP growth rate				−0.021* (0.010)	−0.010 (0.009)
Year 2018		−0.953*** (0.049)	−0.962*** (0.031)	−0.945*** (0.060)	−1.050*** (0.045)
Year 2020		0.602*** (0.048)	0.608*** (0.026)	0.492*** (0.066)	0.562*** (0.047)
City fixed effects	N	N	Y	N	Y
Adjusted R^2	0.010	0.557	0.839	0.523	0.844
No. observations	818	818	818	655	655

Note: This table reports the relationship between Fintech credit risk and real-economy risk. The dependent variable is the 30 days probability of default of Fintech loans in percentage points. Column (II) adds year dummies only (no city fixed effects), while column (III) adds both year dummies and city fixed effects. The same logic applies to columns (IV) and (V), which additionally control for GDP per capita and GDP growth. The 2018 and 2020 year dummies use 2019 as the reference year. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

We now regress the PD of the platform lender against defaults and economic conditions at the city level. Table 2 presents the results. First, there is a clear positive correlation

between AI-based platform credit defaults and default rates in the real economy (column I). Yet this relationship mainly reflects the common upward trend in 2018–2020. To disentangle the roles of the year dummies and the city fixed effects, column (II) adds only the year dummies (using 2019 as the reference year) without city fixed effects, while column (III) layers city fixed effects on top. The sign of “Defaulters” already flips from positive to negative in column (II), so the year dummies alone are enough to account for the sign reversal; adding city fixed effects in column (III) sharpens the estimate but is not what drives the change in sign.

Second, after accounting for time effects and city fixed effects, the relationship between AI-based platform credit risk and real-economy risk becomes negative. An increase of one default per 1,000 population (a 25% increase) is associated with a reduction in the platform PD rate by 0.13 to 0.27 percentage points. The negative relationship suggests that in regions where real economy risk levels are higher, the AI-based lending system takes a stricter, more cautious strategy in controlling credit risk.

Third, columns IV and V show that regions with higher income per capita and faster economic growth exhibit lower credit risk from the platform lender. In Appendix [A](#), we report the analysis based on the 90-day probability of default, which are consistent with those using the 30-day definition.


3.3 AI-based Lending under a Regional Shock: Covid-19

A natural concern about machine-learning credit models is that they are trained on past data and may perform poorly when conditions change abruptly. The Covid-19 shock is such a change. In this subsection we use daily data on the credit portfolio at the city level in the period January 2018–February 2021 to evaluate the effect of risks from Covid-19 shocks. We consider the total loans granted to borrowers in this city as a credit portfolio. On each day, there are new loans disbursed as well as past loans that are delinquent. Then we can approximate the probability of default (PD) by the ratio of delinquent loan amount

to the total outstanding loans. This results in 17,702 observations across 337 cities.

For our analysis, we demarcate the Covid shock based on key events in the policy response. In China, the Covid-19 shock began in Wuhan on 23 January 2020, and soon led to national wide lockdowns. Hubei province (where Wuhan is located) instituted a particularly severe lockdown, which lasted for over three months. On 2 May, Hubei downgraded the emergency response to Covid-19 from the highest level (level 1) to a lighter regime (level 2). By May, most provinces were in the recovery process with no strict restrictions on traffic in and out. Hence, we define the shock period as that from the weeks of 23 January to 4 May 2020.

Figure 5: Probability of default (PD) of AI-based platform credit remained low




Note: The solid vertical lines indicate the week of 23 January and 4 May 2020 (dates in which the lockdown began and ended in Wuhan). The probability of default (PD) rate is defined as the 90 days delinquency rate of SME lending. Source: authors' calculations.

In Figure 5, we compare the PD of SME loans in Wuhan, Beijing and Shanghai. The PD in Wuhan rises first and sharply. Due to the policy of precise prevention and containment, SMEs in Shanghai were less affected compared with the national average. Beijing's lockdown policy was relatively strict, so defaults were also at a relatively high level.

In Figure 6, we plot the change in PD between two periods before and after the shock in different cities. We can see that, not surprisingly, PDs rose more in cities with more

Figure 6: Which cities saw a higher increase in credit risk?



Note: The number of cases refers to total confirmed cases in each city by July 2020, when the total number was only increasing by a single digit number per day. The population of each city is from the 2019 annual statistics. The trade volume with Wuhan of a city is calculated from a sample of a Chinese e-commerce platform's transactions in a period of 2020. Source: authors' calculations.

infected people (left-hand panel). Moreover, using the trade volume of each city with Wuhan in the previous year to represent the economic ties with Wuhan, we can see a weakly positive correlation with the increase in PDs (right-hand panel).

This pattern is confirmed with simple regressions (Table 3). PDs were higher in cities with longer lockdowns and more cases. In particular, the Covid shock increases the PD in 2020 by 0.416 percentage points (column I). The lockdown has a persistent effect, lasting to the end of 2020. One more week of lockdown time resulted in a 0.017 percentage increase in PD, even after the lockdown ended.

In addition to the probability of default, the delinquency data at different frequencies allow us to estimate a loss given default (LGD). Figure 7 shows the LGD based on delinquencies at 180 days vs 30 days past due. The average value is 60.0%—similar to the typical LGD used in standard approaches in banking lending. Based on an alternative definition, based on loans that were delinquent at 90 days, the average LGD is 70.0%.


We also examine the correlation of default rates across cities. Table 4 shows that cities closer to Wuhan experienced higher correlations in default rates during the pandemic.

Table 3: Regressions of PD against Covid shock, lockdown measures and cases

	Dependent variable: probability of default (%)			
	(I)	(II)	(III)	(IV)
Covid dummy	0.416*** (0.001)			
Lockdown time (weeks)		0.037*** (0.001)	0.017*** (0.002)	
Log(number of cases)				0.044*** (0.005)
After lockdown dummy			0.384*** (0.018)	0.574*** (0.015)
City fixed effects	Y	Y	Y	Y
Adjusted R^2	0.095	0.093	0.131	0.135
No. observations	10,917	10,917	10,917	10,917

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure 7: Distribution of loss given default (LGD) across loans



Note: We calibrate the loss given default rate by considering the loans borrowed by SMEs in one city is an asset, and use the ratio of loss after 180 days, to the amount of loans delinquent after 30 days (or 90 days). The average of LGD (30) is 0.600, with median 0.588 and standard deviation 0.187. The average of LGD (90) is 0.700, with median 0.704 and standard deviation 0.177. Source: authors' calculations.

This geographic pattern reflects the spread of both the virus and the economic shock through supply chain and trade linkages.

Table 4: Regressions of correlations against geographic distance

	Dependent variable: $\rho_{i,2020} = \text{correlation}(\text{PD}_{\text{city } i}, \text{PD}_{\text{Wuhan}})_{2020}$			
	(I)	(II)	(III)	(IV)
$\rho_{i,2019}$	0.2822*** (0.0517)	0.2285*** (0.0536)	0.1977*** (0.0568)	0.1977*** (0.0568)
Distance from Wuhan (1,000 km)	-0.1463*** (0.0043)	-0.1325*** (0.0272)	-0.1173*** (0.0283)	-0.1173*** (0.0283)
Lockdown level		0.0029*** (0.0010)		
Cumulative lockdown level			0.0015*** (0.0005)	0.0015*** (0.0005)
Intercept	0.7516*** (0.0336)	0.6220*** (0.0572)	0.2881* (0.1542)	0.2881* (0.1542)
City fixed effects	Y	Y	Y	Y
Adjusted R^2	0.2104	0.2342	0.2381	0.2381
No. observations	334	336	336	336

Note: Source: authors' calculations. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

4 Dynamic Credit Adjustment in the AI-based Lending Model

The previous section showed that credit losses on the AI-based portfolio remained contained, even during the Covid-19 shock. We now turn to how the AI system manages risk after origination. As described in Section 2, the system updates credit admission, lines of credit and interest rates continuously as new data arrive. This is the function of AI that goes beyond screening at origination.


In addition to the limited PDs for online borrowers, the platform lender saw a very small rise in defaults for the (much larger) pool of offline lenders, for whom recourse after default may be much more difficult. Why is this the case? We show that dynamic credit adjustment played an important role.

4.1 Aggregate Dynamic Adjustment During the Covid-19 Shock

Since the AI-based platform credit system is built on real-time big data (primarily payment transaction data), it must continuously process approximately RMB 360 bn (~USD 50 bn)

of transactions per day and respond instantaneously. For each new credit application, the system makes decisions on approval, credit limits, and interest rates. For each existing borrower, the system determines whether to increase or decrease the interest rate, expand or reduce the credit limit, and approve or reject new loan applications. Consequently, the AI-based system must balance at least two key objectives: it must control credit risk and ensure profitability, while at the same time maintaining a smooth borrower experience, e.g., by avoiding seemingly arbitrary increases in interest rates.

Figure 8: Credit admissions depend on PD; removal of access for some lets others enter




Note: The Credit Admission rate is calculated from a sample of three million potential and existing borrowers, indicating the proportion of potential borrowers admitted each day. The PD rate is calculated from a large representative sample of existing borrowers. It is the risk of current outstanding credit balance according to daily repayment activity.

First, we show the adjustments in the credit admission ratio (blue) in response to the change in PDs (red). When PDs rose due to the Covid shock in early 2020, the credit admission ratio dropped quickly by about 3 percentage points, to 94% of all borrowers (Figure 8, left-hand panel). Throughout our period, aggregate credit admission rates were lower on days where the aggregate PD was higher (centre panel). Yet more importantly, the AI-powered lender differentiated between borrowers. Whenever the credit offer to individual borrowers is revoked, a similar number of new borrowers receive an offer for

credit (right-hand panel). It is especially higher-quality borrowers (eg those who have not defaulted in the past) who receive new offers.

Figure 9: Lines of credit (LoC) were adjusted dynamically, expanding for good borrowers when they were reduced for others




Note: The lines of credit are calculated from a sample of three million potential and existing borrowers, indicating the mean of LoC for both potential borrowers and existing borrowers on each day. The PD rate is calculated from a large representative sample of existing borrowers. It is the risk of the current outstanding credit balance according to daily repayment activity.

In Figure 9, we show the dynamic adjustment in the line of credit. The average size of lines of credit was curtailed in March and April 2020, yet this was short-lived, and average lines of credit were soon above their previous levels (left-hand panel). Again, higher-quality borrowers saw their line of credit expanded, and this adjustment happened very quickly. As a result, the correlation between the average size of the line of credit and the overall PD rate is weaker than the correlation between credit admissions and PD (centre panel). As with credit admission, lines of credit were expanded for some borrowers whenever they were reduced for others (right-hand panel).

In Figure 10, we show how interest rates developed. Strikingly, as PDs rose, the average interest rate fell sharply (left-hand panel). There were also changes across the distribution of interest rates, as fewer borrowers faced an interest rate over 20%, and more received an

Figure 10: As the pandemic hit, interest rates were reduced



Note: The interest rate is calculated from a sample of three million potential and existing borrowers, indicating the mean of interest rate for both potential borrowers and existing borrowers on each day. We divide consumers into four groups with similar share at the beginning of the sample period. Group 1 includes borrowers with interest rate less than 12.5%. Group 2 includes those with interest rates less than 15%. Group 3 includes borrowers with rates less than 20%, while the remaining borrowers are in Group 4. SHIBOR refers to the Shanghai interbank offered rate.

offer for loans below 12.5% (right-hand panel). This was all the more notable as central bank policy interest rates and interbank rates declined only slightly during the period.

4.2 Borrower-level Evidence: Lines of Credit and Interest Rates

We now zoom in from aggregate patterns to borrower-level dynamics, examining how the AI system adjusts credit limits and interest rates at the individual firm level during the Covid-19 shock. We draw a random sample of 281,503 *potential* offline borrowers, i.e., SMEs evaluated by the platform as prospective clients during the sample period, whether or not ultimately admitted, and track their credit status monthly from April 2019 to December 2020. We focus on offline borrowers because they account for the majority of the platform's SME client base and span a much wider range of industries, making their credit dynamics more representative of how the AI system responds to a macro shock.

Table 5: Descriptive statistics of the credit adjustment dataset

Variable	Mean	P25	Median	P75	Std. Dev.	N
Is admitted	0.80	1	1	1	0.40	5,527,389
Line of credit (LoC) value (RMB)	28,725	5,000	16,000	45,000	34,653	3,891,670
Loan value (RMB)	5,735	0	0	3,500	16,650	144,107
Change in LoC (RMB)	1,035	0	0	0	11,943	3,591,712
- when LoC increases	10,917	5,000	5,000	10,500	17,275	755,520
- when LoC decreases	-13,116	-16,000	-7,000	-2,800	20,224	345,340
Interest rate (percentage points)	16.01	12.60	16.20	19.80	3.90	3,891,670
Change in interest rate	-0.11	0.00	0.00	0.00	1.34	3,591,712
- when interest increases	3.66	1.80	1.80	5.40	2.57	104,847
- when interest decreases	-3.79	-5.40	-3.60	-1.80	2.61	208,328
Sales value (RMB, monthly)	12,901	1,113	3,508	9,619	112,441	4,505,866
Sales orders (monthly)	163.20	25.00	64.00	166.00	357.54	4,505,866
Gender (Is female)	0.50	0	0	1	0.50	263,209
Age	40.80	34	40	48	9.02	263,209
PD (percentage points)	2.23	1.84	2.08	2.63	0.44	21

Note: Descriptive statistics for the offline-borrower sample used in Sections 4.1. The sample is a balanced borrower-by-month panel of 263,209 borrowers across 21 months (April 2019–December 2020). “Is admitted” is an indicator that equals one in all months from the borrower’s first positive credit month onward. Line of credit, interest rate, and sales variables are observed only for borrower-months with active credit access. Change variables require the borrower to be observed in two consecutive months. PD is the platform-wide 30-day delinquency rate (see Section 3); it is a monthly time series that takes one value per month.

Table 5 reports descriptive statistics. Among admitted borrower-months, the average credit line is RMB 28,725 (~USD 4,100) and the average interest rate is 16.0%. Is admitted equals one in 80% of borrower-months in the balanced panel.³ Adjustments on the two margins differ sharply in frequency. Credit lines are adjusted in 30.6% of borrower-months (21.0% increases, 9.6% decreases), implying a typical borrower sees a credit-line change roughly once every 3.3 months. Interest rates are adjusted far less frequently, in only 8.7% of borrower-months (2.9% up, 5.8% down), or roughly once every 11.5 months per borrower. The platform-wide probability of default ranges from 1.75% to 2.97% across the 21-month window.

³The 80% borrower-month admission rate in Table 5 is not directly comparable with the 94–98% admission rate in Figure 8. The figure is a daily cross-sectional rate computed on a separate three-million-borrower sample: the share of potential borrowers in admitted status on each day. The table is a longitudinal share computed on the 263,209-borrower balanced panel: the 20% of observations with value 0 are pre-entry borrower-months for borrowers who were eventually admitted, not borrowers being rejected.

4.2.1 Dynamic Adjustment in the Line of Credit

To investigate the relationship of adjustment in the line of credit (LoC) and the risk of the aggregate portfolio, we now regress changes in LoC value against the risk level of the previous month:

$$\Delta\text{LoC}_{i,t} = \beta \cdot \text{PD}_{t-1} + \gamma \cdot X_{i,t} + \alpha_i + \varepsilon_{i,t} \quad (1)$$

where PD_{t-1} is the average of the daily PD rate in the previous month, $X_{i,t}$ includes individual time varying variables, and α_i is the individual fixed effect. We run two sets of regressions, first only control PD_{t-1} and then add the time-varying controls $X_{i,t}$.

We use the inverse hyperbolic sine transformation of the change in LoC, $\text{IHS}(\Delta\text{LoC})$, which is approximately linear around zero but behaves like a logarithm for large changes. It allows us to treat both positive and negative changes at once and can be interpreted in terms of semi-elasticity for large changes.

We first analyze the extensive margin of LoC adjustments. In Columns (I)–(III) of Table 6, the dependent variable is an indicator for whether the borrower’s LoC changes in each month. We find that adjustments occur less frequently in higher-risk periods. This pattern remains robust after controlling for time-varying shop operating information and the borrower’s current credit level. The evidence suggests that, overall, the platform adjusts credit lines more cautiously when observing a higher aggregate risk. We then examine the overall magnitude of LoC adjustments in Columns (IV)–(VI). Higher risk is associated with a lower net change in credit lines. Interpreting the IHS specification in semi-elasticity terms, a 0.1 percentage-point increase in risk is associated with an approximately 3.5%–7.5% reduction in credit lines.

Table 7 decomposes LoC adjustments by direction to study asymmetries in the platform’s quantity response. Conditioning on months with LoC increases ($\Delta\text{LoC} > 0$, columns I to III), a higher risk level is associated with larger expansions in credit lines. By contrast, conditioning on months with LoC decreases ($\Delta\text{LoC} < 0$, columns IV to VI), higher risk

Table 6: Credit adjustment for individual borrowers relative to past PD

Dependent variable:	Is LoC changed (0/1)			IHS(Δ LoC)		
	(I)	(II)	(III)	(IV)	(V)	(VI)
PD, previous month (ppt)	-0.144*** (0.001)	-0.144*** (0.001)	-0.143*** (0.001)	-0.753*** (0.007)	-0.678*** (0.007)	-0.345*** (0.008)
Log LoC (value), prev. month			-0.015*** (0.000)			-1.177*** (0.007)
Interest rate (ppt), prev. month			-0.002*** (0.000)			0.054*** (0.002)
Shop controls	N	Y	Y	N	Y	Y
Individual FE	Y	Y	Y	Y	Y	Y
R^2	0.155	0.162	0.163	0.061	0.071	0.106
No. observations	3,382,776	2,932,390	2,932,390	3,382,776	2,932,390	2,932,390

Note: Shop controls include the log sales value (RMB) and log sales orders. Standard errors, reported in parentheses, are clustered at the individual borrower level. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

is associated with deeper contractions. Importantly, the risk sensitivity is substantially stronger on the expansion margin than on the contraction margin.

Table 7: Credit adjustment for individual borrowers is asymmetric

Dependent variable:	IHS(Δ LoC)					
	When LoC increases (Δ LoC > 0)			When LoC decreases (Δ LoC < 0)		
	(I)	(II)	(III)	(IV)	(V)	(VI)
PD, previous month (ppt)	0.388*** (0.003)	0.404*** (0.004)	0.379*** (0.004)	-0.051*** (0.006)	-0.084*** (0.007)	-0.070*** (0.006)
Log LoC (value), prev. month			-0.054*** (0.002)			-1.061*** (0.005)
Interest rate (ppt), prev. month			-0.028*** (0.001)			-0.056*** (0.001)
Shop controls	N	Y	Y	N	Y	Y
Individual FE	Y	Y	Y	Y	Y	Y
R^2	0.487	0.495	0.498	0.670	0.688	0.786
No. observations	691,339	625,065	625,065	326,398	282,082	282,082

Note: Standard errors, reported in parentheses, are clustered at the individual borrower level. Shop controls include the log sales value (RMB) and log sales orders. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

4.2.2 Dynamic Adjustment in Interest Rate

In addition to adjusting the available lending value, the AI-powered platform lender adjusts the interest rate. On average, interest rates decrease by 0.07 percentage points per month across the full sample (Table 5). Unlike credit line adjustments, which are relatively frequent (18.0% increases and 9.1% decreases), interest rate changes are infrequent. Interest rates increased in 2.6% of observations, with an average rise of 3.89 percentage points, and decreased in 4.5% of observations, with an average drop of 3.76 percentage points.

As the platform accumulates more information and behavioral data about borrowers over time, interest rates for most borrowers tend to decline mechanically as relationships mature, as shown in Figure 10. To separate the relationship between interest rate adjustments and the aggregate risk from this underlying downward trend, we therefore rely on an individual-level panel dataset. Specifically, we regress changes in the interest rate value against the risk level of the previous month:

$$\Delta Interest_{i,t} = \beta \cdot PD_{t-1} + \gamma \cdot X_{i,t} + \alpha_i + \varepsilon_{i,t} \quad (2)$$

where PD_{t-1} is the average of the daily PD rate in the previous month, $X_{i,t}$ includes individual time varying variables, and α_i is the individual fixed effect.

Similarly, we first focus on the extensive margin of interest rate adjustment. In columns (I)–(III) of Table 8, the dependent variable is an indicator for whether the interest rate changes in a given month. Higher aggregate risk is associated with less frequent repricing. The effect is persistent when conditioning on time varying shop controls and the current credit profile.

Then we use the monthly change in the interest rate (in percentage points) as the dependent variable to analyze the magnitude of adjustments. As shown in columns IV and V in Table 8, higher aggregate risk is associated with an increase in the net rate change.

However, we condition on the state variable of credit profile, the interest rate, and LoC of individual, the effect is reversed. It reveals a counter-cyclical intent of the AI algorithm. By isolating the algorithm’s marginal response to an aggregate risk signal, the negative coefficient -0.140 suggests that for any given borrower profile, an increase in portfolio risk triggers an intentional downward adjustment in rates.

Table 8: Interest rate adjustments for individual borrowers relative to past PD

Dependent variable:	Is rate changed (0/1)			Δ Interest Rate (ppt)		
	(I)	(II)	(III)	(IV)	(V)	(VI)
PD, previous month (ppt)	-0.021^{***} (0.000)	-0.016^{***} (0.000)	-0.007^{***} (0.000)	0.176^{***} (0.002)	0.172^{***} (0.002)	-0.140^{***} (0.002)
Log LoC (value), prev. month			0.004^{***} (0.000)			-0.063^{***} (0.001)
Interest rate (ppt), prev. month			0.007^{***} (0.000)			-0.230^{***} (0.000)
Shop controls	N	Y	Y	N	Y	Y
Individual FE	Y	Y	Y	Y	Y	Y
R^2	0.105	0.111	0.113	0.062	0.072	0.173
No. observations	3,382,776	2,932,390	2,932,390	3,382,776	2,932,390	2,932,390

Note: Standard errors, reported in parentheses, are clustered at the individual borrower level. Shop controls include the log sales value (RMB) and log sales orders. $*** p < 0.001$, $** p < 0.01$, $* p < 0.05$.

To provide a more granular understanding of interest rate adjustment, we conduct separate analyses for samples experiencing rate increases and decreases. Columns (I)–(II) of Table 9 show that when rates increase, the magnitude of the increase is smaller when aggregate risk level is higher. Further, when rates decrease, the magnitude of the change is also smaller (in absolute value) in higher-risk months (Columns (IV)–(V)). Taken together, these results suggest that, conditional on repricing of interest rates, the AI system adjusts the interest rates more cautiously when aggregate risk is elevated.

4.3 Algorithmic Bias in Credit Adjustment

A key concern with using algorithms in lending is whether they introduce bias based on personal characteristics. Since the adjustments documented above are made by the AI

Table 9: Interest rate adjustments are more conservative when risk is high

Dependent variable:	Δ Interest Rate (ppt)					
Sample	When interest increases			When interest decreases		
	(I)	(II)	(III)	(IV)	(V)	(VI)
PD, previous month (ppt)	−0.907*** (0.030)	−0.975*** (0.034)	−0.372*** (0.028)	1.132*** (0.018)	1.187*** (0.020)	−0.306*** (0.020)
Log LoC (value), prev. month			−0.140*** (0.013)			−0.096*** (0.010)
Interest rate (ppt), prev. month			−0.568*** (0.008)			−0.513*** (0.004)
Shop controls	N	Y	Y	N	Y	Y
Individual FE	Y	Y	Y	Y	Y	Y
R^2	0.860	0.875	0.921	0.810	0.820	0.869
No. observations	101,367	85,084	85,084	194,965	173,144	173,144

Note: Standard errors, reported in parentheses, are clustered at the individual borrower level. Shop controls include the log sales value (RMB) and log sales orders. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

system without human intervention, we examine whether the system treats borrowers differently based on gender, education, and age.

In Table 10, we check the difference in credit adjustment for male and female borrowers. Neither the adjustment in LoC amount nor the percentage rate shows any bias toward female borrowers. However, the male borrowers experience a more significant adjustment to the PD change.

Table 11 shows that the response to different education groups is not linear. While it is decreasing from the junior college group, college, to graduate degree group, the algorithm did not produce biased results for individuals with the lowest educational group (high school or below).

In Table 12 we report the results against age groups. It shows that credit line adjustments are not biased against younger borrowers (aged 25 and under), with the largest adjustment in LoC for the eldest age group.

Table 10: Credit adjustment and gender

	Dependent variable: Change in credit line			
	Amount		Percentage	
	(I) Male	(II) Female	(III) Male	(IV) Female
PD, prev. month avg. (ppt)	−1,167.5*** (16.4)	−1,082.2*** (14.5)	−8.588*** (0.143)	−8.344*** (0.133)
Sales (amount)	−0.000 (0.000)	−0.001*** (0.000)	−0.000 (0.000)	−0.000** (0.000)
Sales (orders)	0.771*** (0.034)	0.780*** (0.035)	0.003*** (0.000)	0.003*** (0.000)
LoC (amount)	0.174*** (0.000)	0.181*** (0.000)	0.001*** (0.000)	0.001*** (0.000)
Individual FE	Y	Y	Y	Y
R^2	0.163	0.164	0.104	0.107
No. observations	1,460,624	1,471,766	1,444,291	1,457,229

Note: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

Table 11: Credit adjustment for individual borrowers relative to past PD by education

	Dependent variable: Change in credit line (amount)			
	High school or below	Junior college	College	Graduate Degree
PD, prev. month avg. (ppt)	−1,121.3*** (11.0)	−1,690.1*** (157.3)	−1,206.4*** (125.5)	−607.4* (261.4)
Sales (amount)	−0.000** (0.000)	0.001 (0.001)	−0.000 (0.001)	−0.001* (0.001)
Sales (orders)	0.778*** (0.024)	0.553+ (0.302)	0.804** (0.309)	0.392 (0.402)
LoC (amount)	0.178*** (0.000)	0.149*** (0.003)	0.156*** (0.002)	0.152*** (0.006)
Individual FE	Y	Y	Y	Y
R^2	0.164	0.167	0.170	0.158
No. observations	2,864,142	24,315	36,470	7,463

Note: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, + $p < 0.1$.

Table 12: Credit adjustment for individual borrowers relative to past PD by age group

	Dependent variable: Change in credit line (value, RMB)				
	Age ≤ 25	26–35	36–45	46–55	> 55
PD, prev. month avg. (ppt)	−154.5* (76.9)	−1,031.5*** (22.3)	−954.0*** (18.5)	−1,017.9*** (18.8)	−1,087.7*** (42.8)
Sales (amount)	0.00 (0.00)	−0.00 (0.00)	−0.00 (0.00)	−0.00 (0.00)	0.00 (0.00)
Sales (orders)	0.01 (0.18)	0.95*** (0.05)	0.84*** (0.04)	0.60*** (0.04)	1.19*** (0.12)
LoC (amount)	0.18*** (0.00)	0.14*** (0.00)	0.14*** (0.00)	0.16*** (0.00)	0.18*** (0.00)
Individual FE	Y	Y	Y	Y	Y
R^2	0.226	0.148	0.143	0.149	0.160
No. observations	43,494	872,372	1,053,586	808,265	154,673

Note: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

5 The Limits of AI: Platform Participation and Incentives to Repay

The dynamic adjustments documented in the previous section describe one half of the platform’s credit technology: the lender’s ability to re-price and re-size exposures in real time. This is what AI does well: it addresses the lender’s information problem. Yet, even the most accurate prediction and the most responsive limit-setting cannot prevent default on loans already outstanding, and these loans are entirely uncollateralised. AI can predict who is likely to default, but it cannot by itself make a borrower repay. Why did borrowers not choose to strategically default on their loans? In this section, we propose and test the hypothesis that borrowers repay because default would sever their access to the platform’s commercial ecosystem. The value of platform participation thus serves as implicit collateral, sustaining repayment incentives in the absence of conventional collateral.

Testing this hypothesis requires data on both credit outcomes and the borrower’s stake in the platform’s ecosystem. We therefore focus on online SME borrowers who operate e-commerce shops on an associated platform, as they are the set of borrowers for whom platform participation is directly observable and economically meaningful.

5.1 Conceptual Framework

To motivate our empirical strategy, we develop a simple framework for a platform borrower’s default decision. Consider a borrower with an outstanding loan obligation L who earns platform revenue R in the current period. The borrower chooses between repayment and default. If the borrower repays, she retains access to the platform and receives a discounted continuation value δV , where V represents the expected future surplus from platform participation. If the borrower defaults, she keeps current-period cash flow but is

excluded from the platform ecosystem, forfeiting V . The borrower's payoffs are:

$$\text{Repay: } R - L + \delta V$$

$$\text{Default: } R$$

The borrower defaults if and only if $\delta V < L$. That is, default occurs when the discounted value of continued platform access falls below the loan obligation. The key insight is that the default decision depends on the platform continuation value V , not merely on current cash flow R .

We decompose the continuation value as $V = V_0 + \alpha S$, where V_0 is the baseline value of platform access and S represents the borrower's accumulated platform-specific investments, such as pay-for-performance (P4P) advertising expenditures, reputation capital (seller star ratings), and customer relationships built within the platform. These investments are sunk and non-transferable: they have value only within the platform ecosystem, and a borrower who defaults and is excluded from the platform loses the entire return on them. Higher platform revenue R can therefore reduce default through two distinct channels: an *income channel*, by relaxing the current-period cash-flow constraint, and an *incentive channel*, by strengthening the borrower's stake in preserving V . Because the two channels are observationally equivalent under a single correlation of R with default, we need a research design that induces variation in R independent of S (Prediction 1), and tests for differential responses across borrowers with different S (Predictions 2–3).

This framework generates three testable predictions:

1. *Platform dependency reduces default*: Higher platform revenue R is associated with greater platform engagement and a higher continuation value V , reducing the likelihood of default. This prediction is consistent with either channel.
2. *Platform-specific investment amplifies the effect* (incentive channel): The effect of platform revenue on default should be stronger for borrowers with larger platform-

specific sunk costs S , since these borrowers have a higher continuation value V and thus more to lose from exclusion. Formally, $\partial^2 \Pr(\text{Default}) / \partial R \partial S < 0$.

3. *No amplification under pure income channel*: If the sales–default link operates solely through cash flow (ability to pay), then the effect should not vary with platform-specific investment S , since cash-flow relief from higher revenue does not depend on the borrower’s sunk costs.

We test Prediction 1 by instrumenting shop revenue with a platform-specific certification (the Gold Seller label); Predictions 2 and 3 are then tested via heterogeneity in the response by P4P advertising and seller-star rating, our two proxies for platform-specific investment.

5.2 Data and Sample

We start with a random sample of 499,100 borrowers with at least one loan originated between January 1, 2018, and March 31, 2021, and construct a borrower-month panel running from each loan’s disbursement month to its end month (settlement if repaid, delinquency if defaulted, or March 2021 if still outstanding). The default dummy equals 1 in the month any of the borrower’s loans transitions into default, and 0 otherwise; for borrowers with multiple outstanding loans, we count active loans and sum amounts and balances. The panel has 6.5 million borrower-months, with a median of 12 active months per borrower.

We link these records to borrower information from the platform lender and shop-operation data from the e-commerce platform, the latter covering roughly two-thirds of borrower-months, since some online sellers are not actively operating their shops on the platform in every month, or have exited the platform altogether. After dropping observations missing any control, the analysis sample contains 2,299,015 borrower-months on 367,173 borrowers, with a monthly default-event rate of 0.2 pp.

Table 13 reports the summary statistics for our analysis sample. On a typical borrower-

Table 13: Descriptive statistics at borrower-month level

Variable	Mean	P25	Median	P75	Std. Dev.	N
<i>Loan information</i>						
Default dummy (0/1)	0.002	0.000	0.000	0.000	0.046	2,299,015
Default amount share (%)	0.21	0.00	0.00	0.00	4.58	2,299,015
Loan amount (yuan)	78,710	12,000	33,000	87,000	122,728	2,299,015
# of Loans	3.67	1.00	2.00	4.00	4.56	2,299,015
Outstanding balance (yuan)	3,235	0	0	0	20,129	2,299,015
Interest Rate (%)	14.76	13.87	14.96	15.81	2.22	2,299,015
<i>Shop information</i>						
Sales (yuan)	59,726	3,024	12,848	47,439	139,440	2,298,986
# of Buyers	354	16.0	66.0	261	866	2,266,116
Gold seller dummy (0/1)	0.163	0.000	0.000	0.000	0.369	2,299,015
Gold seller (lag 1)	0.17	0.00	0.00	0.00	0.37	1,040,374
Gold seller (lag 2)	0.17	0.00	0.00	0.00	0.38	936,512
Deposit amount (yuan)	1,065	0	1,000	1,001	1,502	2,299,015
Shop open years	6.08	3.80	5.72	8.10	3.02	2,299,015
Shop rating (0 to 1)	0.989	0.986	0.993	0.997	0.017	2,299,015
Customer service online time (hour)	10.4	4.18	7.79	12.0	10.8	2,298,971
P4P participant (0/1)	0.827	1.000	1.000	1.000	0.378	2,299,015
<i>Borrower Characteristics</i>						
Age	35.7	29.0	33.0	39.0	9.79	367,173
Years used digital payment	7.19	5.00	7.00	9.00	3.04	367,173
Is property owner	0.908	1.000	1.000	1.000	0.289	367,173
Is car owner	0.586	0.000	1.000	1.000	0.493	367,173

Note: Borrower-month statistics on the analysis sample (367,173 unique borrowers, 2018–2021). To mitigate the influence of a small number of extreme observations on the mean and standard deviation, the right-skewed continuous variables (sales, number of buyers, loan amount, number of loans, customer service online time, and deposit) are winsorized at the 99th percentile. Percentiles are unaffected by the winsorization because the cap equals the 99th percentile itself.

month, the average loan size is RMB 78,710, the interest rate is 14.76%, and a borrower holds 3.7 active loans with an average outstanding balance of RMB 3,235. The monthly default event rate is 0.2%, and the corresponding share of outstanding balance in default is 0.21%; across the full sample of completed loans, the ever-default rate is 2.05%. At the shop level, the average monthly sales revenue is RMB 59,726 from 354 buyers, and the seller is active on the platform's customer-service software for 10.4 hours per month on average. Shops in the sample have operated on the platform for a median of 5.7 years, with an average customer-service rating of 0.989. To limit the influence of a small num-

ber of extreme observations, the right-skewed continuous variables (sales, buyers, loan amount, number of loans, customer-service time, and deposit) are winsorized at the 99th percentile.

Platform-specific investments are widespread: about 82.7% of borrower-months feature active participation in the platform’s pay-for-performance (P4P) advertising program, giving us variation in platform-specific capital with which to test the incentive channel in Section 5.6. Because the platform and the lender are tightly linked, default carries not only the direct financial consequences but also exclusion from the ecosystem that hosts the borrower’s operations, reputation, and advertising investments. This setting enables us to measure the degree of ecosystem dependence and assess its causal role in default risk in AI-based platform lending.


5.3 Identification: The Gold Seller Instrument

We proxy a borrower’s platform dependence by shop-level revenue. The empirical challenge is endogeneity: successful sellers have both higher sales and characteristics, such as better product quality, stronger management, and deeper customer relationships, that independently affect default. We address this by instrumenting sales with a platform-specific certification, the “Gold Seller” label, whose assignment is governed by an algorithm internal to the e-commerce platform and is not observed by the lender at credit-setting time.

The Gold Seller label is a monthly qualification granted by the e-commerce platform to selected sellers in each product category, based on shop operations (sales rank, operating length, number of orders, customer visits) and service quality (customer ratings, service ratings, dispute incidence). Certified sellers receive a visible badge on their shop and product pages that raises consumer trust and, mechanically, lifts sales in the current and following one to two months.

Although higher-revenue sellers are more likely to be certified, the label is available

Figure 11: Identification Strategy: Gold Seller as Instrumental Variable



Note: The e-commerce platform certifies Gold Sellers monthly based on transaction volume and service quality. Gold Seller status boosts traffic and sales (first stage). The fintech lender provides payment services to the platform but does not have access to Gold Seller certification data, ruling out a direct channel through loan terms (exclusion restriction).

to sellers at every part of the sales distribution: the certification rate rises roughly monotonically from 5.2% in the bottom sales quintile to 93.0% at the top, but exceeds zero in every quintile.⁴ In any given month, 16.3% of sellers are certified, and 52% of sellers in our sample receive the label at least once over 2018–2021.

Relevance is immediate: certification produces a sharp, short-lived increase in sales. The identifying assumption is exogeneity, namely that Gold Seller status affects default only through its effect on sales, and we argue it from two institutional features.

- (a) *Organizational separation.* The Gold Seller algorithm is operated by the e-commerce arm and uses only shop-operation and customer-service inputs, none of which enter the platform lender’s credit-scoring model.

⁴The certification rates across the five sales quintiles (lowest to highest) over the full sample period are 5.2%, 27.3%, 57.7%, 79.8%, and 93.0%.

- (b) *Timing separation.* The certificate is issued on the first day of each month and is not fed into the lender's real-time credit-adjustment system as an input to admission, line size, or interest-rate decisions. The lender therefore cannot condition credit terms on Gold Seller status, a falsification we verify directly in the next subsection.

Together, these two features rule out the most natural direct channel, namely lender response to certification, through which the instrument could otherwise affect default.

We estimate the causal effect of platform dependence on default by two-stage least squares. The first stage instruments log revenue with the Gold Seller dummy and up to two of its monthly lags:

$$Y_{i,t} = \sum_{k=0}^K \delta_k \cdot \text{GoldSeller}_{i,t-k} + \gamma_1 \cdot L_{i,t} + \gamma_2 \cdot S_{i,t} + \alpha_j + \mu_{ct} + \varepsilon_{i,t}, \quad (3)$$

where $L_{i,t}$ collects loan-level controls (loan amount, number of loans, outstanding balance, interest rate), $S_{i,t}$ collects shop- and borrower-level controls, α_j is industry fixed effect, and μ_{ct} is city-by-month fixed effect. We report estimates using $K = 0, 1, 2$ to trace the sensitivity of the results to instrument count; the just-identified $K = 0$ specification has the strongest first stage, while the over-identified $K = 1, 2$ specifications admit a Sargan check of joint instrument validity.

The second stage uses the first-stage fitted values to identify the causal effect on default:

$$\text{Default}_{i,t} = \beta \cdot \hat{Y}_{i,t} + \gamma_1 \cdot L_{i,t} + \gamma_2 \cdot S_{i,t} + \alpha_j + \mu_{ct} + \varepsilon_{i,t}. \quad (4)$$

Standard errors are clustered at the borrower level. We also report the corresponding OLS specification for comparison:

$$\text{Default}_{i,t} = \gamma_0 + \gamma_1 Y_{i,t} + \gamma_2 \text{GoldSeller}_{i,t} + \gamma_3 L_{i,t} + \gamma_4 S_{i,t} + \alpha_j + \mu_{ct} + u_{i,t}. \quad (5)$$

5.4 Platform Dependency and Default: Baseline IV Estimates

Table 14 reports the main IV estimates. Columns (1)–(3) use a monthly default dummy (scaled to 0–100) as the outcome; columns (4)–(6) use the share of outstanding balance in default (also in percent). Within each outcome, the three columns progressively add Gold Seller lags to the instrument set: column 1 uses only the current-month dummy (just-identified); columns 2 and 3 add the one- and two-month lags (over-identified). The coefficient of interest in Panel B is log sales, our proxy for platform dependency.

First stage. Panel A shows that Gold Seller certification has a sharp, quantitatively large, and highly significant effect on log sales. In the just-identified specification (column 1), receiving the certificate is associated with a 2.13 log-point increase in current-month sales. The first-stage F -statistic exceeds 480,000, ruling out any weak-instrument concern. Adding the one- and two-month lags preserves the significance of the current-month instrument (coefficients 1.37–1.55) and produces individually significant lagged effects of comparable magnitude (0.65–1.07), confirming that the sales boost from certification persists for two to three months before decaying.

Baseline 2SLS. Panel B shows that platform dependency causally reduces default. In the headline specification (column 1), a one-log-unit increase in instrumented sales reduces the monthly probability of default by 0.068 pp ($SE = 0.003$). The estimate is virtually unchanged when we add the first lag (-0.064^{***} , column 2) and slightly larger when both lags are added (-0.075^{***} , column 3). Using the default amount share as the outcome yields an equally robust picture (columns 4–6): log sales enters with a coefficient of -0.049^{***} to -0.059^{***} , roughly four-fifths of the default-dummy estimate.

Panel C reports the corresponding OLS estimates. The OLS coefficients on log sales (-0.057 to -0.068) are smaller in absolute magnitude than the 2SLS estimates across every column and outcome, with the Wu-Hausman test rejecting exogeneity of log sales in

the just-identified specification ($p = 0.017$). This pattern, 2SLS larger than OLS, is consistent with the IV identifying a local average treatment effect among platform-dependent compliers for whom the sales–default link is strongest, rather than a bias from omitted seller quality, which would push in the opposite direction; selection of this kind is known to shape estimated credit-market relationships more broadly (Yannelis and Zhang, 2023). The direct OLS association between Gold Seller status and default (after conditioning on log sales) is a small -0.023^{**} (column 1), further evidence that the Gold Seller–default association is mediated almost entirely through sales.

Economic Significance. A 10% increase in instrumented sales lowers monthly default by 0.0068 pp (95% CI $[-0.0074, -0.0062]$). The relevant baseline is Against the baseline of the complier default rate 0.25 pp (Appendix C), a 10% sales increase delivers a 2.7% (2.5–3.0%) relative reduction in default. The first-stage coefficient of 2.13 log points is essentially one standard deviation of log sales ($\sigma = 2.21$), so a per-certification and a per-SD exercise coincide; the former is the policy-implementable unit.

Lifetime magnitude. The lifetime effect comes directly from the loan-level IV in Appendix C.3. Multiplying the monthly coefficient by the loan’s maturity would misstate it: cumulative default is $1 - (1 - p)^T$, and the event study shows the first stage decays within two to three months. Ever being Gold Seller during the loan’s life lowers lifetime default by 0.33 pp, or 15% of the 2.20% baseline. A platform-wide 10% sales increase lowers it by 0.047 pp, or 2.1% of baseline.

Translation into expected loss. The loss given default is 60% (Section 3.3). Combined with a lifetime default probability of 2.20% from the loan-level sample (Appendix C.3), the baseline expected loss on the mean loan of RMB 40,117 is RMB 530. A Gold Seller certification lowers this by RMB 80; a platform-wide 10% sales increase lowers it by RMB 11. The default-amount share yields a coefficient about four-fifths as large as the default dummy, so the RMB magnitudes are not driven by a drop in small-balance delinquencies alone.

Table 14: Platform dependency reduces default

	Default dummy ($\times 100$)			Default amount share (%)		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: First Stage for Log (Sales)</i>						
Gold seller dummy (0/1)	2.133*** (0.005)	1.550*** (0.005)	1.365*** (0.007)	2.133*** (0.005)	1.550*** (0.005)	1.365*** (0.007)
Gold seller (lag 1)		1.072*** (0.005)	0.829*** (0.006)		1.072*** (0.005)	0.829*** (0.006)
Gold seller (lag 2)			0.645*** (0.007)			0.645*** (0.007)
F-Statistic	487,898	134,401	42,759	487,898	134,401	42,759
Adjust R^2	0.400	0.432	0.450	0.400	0.432	0.450
Sargan p -value	—	0.021	0.060	—	0.009	0.098
Wu-Hausman p -value	0.017	0.342	0.493	0.101	0.683	0.520
<i>Panel B: Two-Stage Least Squares</i>						
Log (sales)	-0.068*** (0.003)	-0.064*** (0.005)	-0.075*** (0.008)	-0.053*** (0.003)	-0.049*** (0.004)	-0.059*** (0.006)
Log (loan amount)	0.007* (0.003)	-0.001 (0.004)	-0.002 (0.007)	0.002 (0.002)	-0.006 (0.004)	-0.008 (0.006)
Number of loans	0.015*** (0.001)	0.014*** (0.001)	0.013*** (0.002)	0.009*** (0.001)	0.008*** (0.001)	0.008*** (0.002)
Shop open years	-0.005** (0.002)	-0.004 (0.003)	-0.006 (0.004)	-0.004* (0.002)	-0.004 (0.002)	-0.005 (0.004)
<i>Panel C: Ordinary Least Squares</i>						
Log (sales)	-0.057*** (0.002)	-0.058*** (0.004)	-0.068*** (0.006)	-0.046*** (0.002)	-0.046*** (0.003)	-0.053*** (0.005)
Gold seller dummy (0/1)	-0.023** (0.008)	-0.040*** (0.012)	-0.051** (0.019)	-0.014* (0.006)	-0.034** (0.011)	-0.042* (0.017)
Gold seller (lag 1)		0.027* (0.013)	-0.009 (0.019)		0.031** (0.012)	-0.005 (0.017)
Gold seller (lag 2)			0.049* (0.021)			0.038* (0.018)
Loan Controls	YES	YES	YES	YES	YES	YES
Shop and Borrower Controls	YES	YES	YES	YES	YES	YES
Industry FE	YES	YES	YES	YES	YES	YES
City-Month FE	YES	YES	YES	YES	YES	YES
Observations	2,266,337	1,025,290	462,275	2,266,337	1,025,290	462,275

Note: 2SLS and OLS regressions on the borrower-month panel. Outcome in cols (1)-(3) is the default dummy multiplied by 100; in cols (4)-(6) it is the share of the borrower's total outstanding balance that is in default (in %). Columns vary in the number of Gold Seller instruments used. Standard errors clustered at the borrower level. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, + $p < 0.10$.

Instrument diagnostics. The lower panel of Table 14 reports standard IV diagnostics. The Wu-Hausman test rejects exogeneity of log sales for the just-identified specification ($p = 0.017$ for the default dummy) but is less powered in the over-identified columns. For the over-identified specifications, the Sargan test does not reject joint instrument validity at conventional levels in column 3 ($p = 0.060$); the borderline rejection in column 2 ($p = 0.021$) reflects the inclusion of a single lag that partially picks up anticipation effects, and motivates reporting column 1 as the headline estimate.


Robustness and extensions (appendix). Three additional exercises in Appendix C support the baseline estimate. First, Table A8 substitutes three-month cumulative sales for the one-month measure and yields qualitatively identical conclusions, ruling out a short-horizon artefact. Second, an Abadie (2003) κ -weighted complier profile shows that the IV-LATE applies to a marginal, platform-dependent, higher-risk subgroup: the complier share is $\hat{\pi}_c = 0.54$, and compliers have 10% lower sales, 16% fewer buyers, and a 24% higher baseline default rate than the full sample, exactly the sellers for whom the platform-incentive channel would be expected to bind most strongly. Third, re-estimating at the loan level, with each observation one completed loan and the instrument the fraction of loan-life months during which the borrower was a Gold Seller, preserves the direction and significance of the effect and shows that it is concentrated among loans of longer life and among P4P participants.

5.5 Validating the Instrument

The exclusion restriction identified in Section 5 rests on the claim that Gold Seller status affects default only through sales. We test it directly along three fronts: the dynamics of the first stage and reduced form, a falsification on the lender-set loan terms, and a set of auxiliary exclusion checks reported in the appendix.

Event-study dynamics. Figure 12 reports a Sun–Abraham (2021) event study around each seller’s first Gold Seller certification, using a common pool of never-treated controls. Panel (a) (first stage) shows that log sales jumps sharply at $k = 0$ and decays over the following two to three months, precisely the dynamic we need the instrument to have. Panel (b) (reduced form) shows a mirror-image contemporaneous *decrease* in default at $k = 0$ and $k = +1$, whose magnitude scales to the first-stage bite. The coefficient path before $k = 0$ is flat in both panels, supporting the identifying assumption that sellers are not on a pre-trend at the moment of certification. Appendix C reports cohort-specific versions of this event study (for sellers first certified in 2018, 2019, and 2020) and a horizon-by-horizon Wald-IV ratio that is stable around $k = 0$ and becomes unstable only after the first stage has decayed, both consistent with the exclusion restriction holding at the contemporaneous horizon that drives the main estimate.

Figure 12: Event study around the first Gold Seller certification



Note: Sun–Abraham (2021) interaction-weighted event-study coefficients around each seller’s first Gold Seller certification. Estimation uses borrower (`short_id`) and city×month fixed effects. Bars show 95% confidence intervals based on standard errors clustered at the borrower level. Baseline period $k = -1$ is omitted by construction. Cohort-specific event studies and the horizon-by-horizon Wald-IV ratio are reported in Appendix C.

Falsification on lender-set loan terms. The identification argument relies on the timing separation between the e-commerce platform (which issues Gold Seller certifications on day 1 of each month) and the platform lender (which does not observe these certifications when setting credit terms). If this assumption holds, Gold Seller status should *not* predict within-borrower changes in the lender-set credit line or interest rate. Table 15 reports the test: within borrowers, changes in the log credit line and interest rate, both in the month of certification and one month ahead, are statistically indistinguishable from zero (point estimates -0.004 and -0.009 for log credit line, 0.007 and -0.002 in percentage points for interest rate). This direct falsification rules out the most natural threat to the exclusion restriction: that the lender adjusts loan terms in response to observed certifications and that those adjustments affect default.

Table 15: Falsification: Gold Seller on lender-set credit line and interest rate

	Log(credit line)			Interest rate (pp)		
	(1) Level	(2) Δ this mo.	(3) Δ next mo.	(4) Level	(5) Δ this mo.	(6) Δ next mo.
Gold seller dummy	-0.102^* (0.040)	-0.004 (0.020)	-0.009 (0.021)	0.045 (0.070)	0.007 (0.016)	-0.002 (0.015)
Controls	YES	YES	YES	YES	YES	YES
Industry + City-Month FE	YES	YES	YES	YES	YES	YES
Observations	47,013	28,712	29,331	47,013	28,712	29,331

Note: Test whether Gold Seller status predicts *lender-set* terms. Interest rate is expressed in percentage points. Δ this month is $x_{i,t} - x_{i,t-1}$; Δ next month is $x_{i,t+1} - x_{i,t}$, testing whether the lender responds to current-month Gold Seller status by adjusting next-month terms. Controls: loan amount and count, shop tenure, online time, good-rate, bail amount, borrower demographics (age, digital age, education, car/house ownership). Standard errors clustered at the borrower level. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, + $p < 0.10$.

Auxiliary exclusion checks. Two additional tests reported in Appendix C reinforce the validity argument. The first asks whether sales actually absorbs the Gold Seller–default association, as the exclusion restriction requires: if Gold Seller affects default only through sales, the reduced-form coefficient should shrink toward zero once log sales is controlled

for. Table A9 confirms this: the Gold Seller coefficient is -0.172^{***} with only fixed effects, shrinks to -0.145^{***} after adding loan and shop controls, and collapses to -0.023^{**} once log sales is added—an 84% attenuation, with no further meaningful reduction from log buyers or log orders.

The second test addresses the concern that Gold Seller status proxies for a time-invariant quality channel rather than the contemporaneous sales mechanism. If quality drives the result, future Gold Seller status should predict current default even after controlling for current sales. Table A10 reports the lead-placebo: the one-month-ahead Gold Seller dummy enters negatively on its own (-0.097^{***}) but attenuates by roughly 40% once log sales is controlled for, with the remaining residual reflecting the seller-improvement trajectory visible in the pre-event portion of the event study rather than a time-invariant quality channel.

We also verify that the platform-dependency measure itself is not driving the result. Table A6 in Appendix C replaces log sales with log number of buyers, which is the channel through which the Gold Seller badge mechanically operates (higher visibility $\rightarrow$ more visits $\rightarrow$ more purchases). The 2SLS coefficients on log buyers are of the same sign and comparable magnitude (-0.082 to -0.096 on the default dummy), indicating that the effect is driven by platform engagement rather than by the particular choice of sales as the proxy.

5.6 Mechanisms: Income Channel vs. Platform Incentives

Having established that platform dependence causally reduces default, we now ask *why*. Two complementary channels can generate this relationship. The *income channel* posits that higher sales simply relax the current-period cash-flow constraint, improving ability to repay.⁵ The *incentive channel*, following our framework in Section 5, posits that sell-

⁵Consistent with this margin, Dobbie and Skiba (2013) find that larger payday loans reduce borrower default.

ers with greater platform-specific capital S have higher continuation value $V = V_0 + \alpha S$ and therefore stronger willingness to repay to avoid exclusion. Predictions 2 and 3 of the framework discipline these channels against the data: if the incentive channel is operative, the estimated effect of sales on default should be stronger for borrowers with larger platform-specific investments; if only the income channel operates, no such differential should appear.

Table 16 tests these predictions. Panel A splits the sample by two proxies for platform-specific investment. Panel B then runs a cash-flow control test that disciplines the income channel directly.

Platform-specific investment amplifies the effect. The first row of Panel A splits the sample by participation in the platform’s pay-for-performance (P4P) advertising program, which requires the seller to make a direct monetary commitment to platform-specific marketing. Among non-P4P sellers, the 2SLS estimate is -0.036^{***} ; among P4P participants, it is -0.074^{***} , roughly twice as large. The second row splits by the seller star rating system, which reflects cumulative transaction volume and reputation accumulated within the platform. Here the contrast is even sharper: among sellers with only star or diamond ratings, the coefficient is -0.018 and statistically indistinguishable from zero, while among crown-rated or higher sellers it is -0.065^{***} . In other words, the sales–default link is *concentrated* among sellers who have made substantial platform-specific investments, and essentially absent among those who have not. This is exactly the pattern predicted by the incentive channel (Prediction 2) and inconsistent with a pure income channel (Prediction 3 would imply no amplification). Further heterogeneity by shop tenure and the owner’s digital age (the owner’s tenure on the platform as a consumer, prior to opening a shop), proxies for general business experience rather than platform-specific capital, is reported in Appendix C and shows no comparable amplification, supporting the interpretation that the P4P and star-rating patterns are capturing platform-specific rather than generic borrower

quality.

Cash flow only partially mediates the effect. Panel B adds log outstanding loan balance, a proxy for current liquidity pressure, to the main IV specification. The 2SLS coefficient on log sales drops from -0.064^{***} (baseline) to -0.026^{***} , a 59% attenuation. Read at face value, this suggests that liquidity accounts for roughly 59% of the sales–default relationship and non-liquidity channels for the remaining 41%. We caveat that loan balance is itself endogenous, so this partition is suggestive rather than definitive; the direction of the implied non-cash-flow residual is nonetheless robustly consistent with the heterogeneity evidence in Panel A.

Interpretation: the two channels are complementary, not substitutes. The model in Section 5 makes it clear that the two channels are not independent: because $V = V_0 + \alpha S$ increases with sales, a seller with high current revenue *both* has an easier cash-flow constraint (income) *and* has more to lose from exclusion (incentive). The heterogeneity results identify the marginal difference in sensitivity across sellers with different S , and this marginal difference is large: the effect is roughly twice as large for P4P sellers as for non-P4P sellers, and effectively zero among sellers who have not accumulated platform-specific reputation. Combined with the partial cash-flow mediation in Panel B, the evidence points to platform-specific capital as an important and distinct mechanism operating on top of the income channel rather than in place of it.

Taken together, the heterogeneity and cash-flow evidence favors the platform-incentive interpretation of our baseline result. The effect of sales on default is substantially stronger, indeed only clearly present, for sellers whose platform-specific investments in advertising and reputation make exclusion more costly. This finding is consistent with Prediction 2 of our framework and supports the hypothesis that the platform ecosystem itself acts as an enforcement mechanism in unsecured lending, complementing the dynamic credit-adjustment technology documented in Section 4.

Table 16: Heterogeneity in Platform Dependency Effect: Income vs. Incentive Channels

	Dep. var.: Default dummy ($\times 100$)	
	(1)	(2)
<i>Panel A: Platform Investment Proxies</i>		
<i>P4P advertising</i>	Non-P4P	P4P
Log (sales)	−0.036*** (0.009)	−0.074*** (0.006)
First-stage F-stats	[15, 377]	[114, 383]
N	169,587	853,856
<i>Seller star rating</i>	Stars/Diamonds	Crowns+
Log (sales)	−0.018 (0.048)	−0.065*** (0.005)
First-stage F-stats	[204]	[133, 295]
N	4,978	1,019,721
<i>Panel B: Cash Flow Control</i>		
	Baseline	+ Loan Bal.
Log (sales)	−0.064*** (0.005)	−0.026*** (0.005)
Log (loan bal.)		0.387*** (0.009)
First-stage F-stats	[134, 401]	[133, 836]
N	1,025,290	1,025,290
Controls & FEs	Yes	

Note: 2SLS estimates of $\log(\text{sales})$ on default (0–100) using two instruments (current and one-month lag Gold Seller) on the borrower-month panel. Panel A splits by the italicized variable. Panel B adds log outstanding balance as a cash-flow control. All use full controls, industry FE, and city $\times$ month FE. First-stage F-stats in brackets. Clustered SEs at the borrower level. Heterogeneity by shop tenure and platform digital age is reported in Appendix C (Table A18). *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, + $p < 0.10$.

6 Conclusion

Artificial intelligence continues to develop rapidly, and is aiding the emergence of new credit intermediaries. AI-powered platform lenders are becoming increasingly important as lenders to small businesses in China, and similar business models are being replicated in a number of other countries around the world. Understanding the unique features of AI-based platform credit and credit risk is thus important for researchers and policymakers.

Using a unique, proprietary dataset from a Chinese platform lender, this paper has shown important insights by which AI-powered lenders can mitigate credit risk and what role they played during the Covid-19 shock. First, AI serves as a real-time risk management tool: even for offline firms that do not depend as much on the value of the platform ecosystem, the dynamic adjustment of credit limits and loan terms helped the AI-powered lender to contain credit losses. Second, AI has a limit: prediction does not enforce repayment. The reliance of online firms on the platform enticed them to repay to maintain the value of platform participation. AI-powered lending is therefore most effective when algorithmic risk management is integrated with platform access and activity.

Our findings contribute to the rapidly growing literature on AI in finance ([Foucault et al., 2025](#)). As AI transforms credit markets from screening and origination to dynamic risk management and enforcement, understanding how these algorithms interact with platform ecosystems becomes increasingly important. Going forward, further research will be needed to assess the quantitative importance of these mechanisms and the extent to which they can be replicated by other lenders aside from large platform firms. It will also be important to assess their importance for growing AI-powered platform activities in other economies around the world.

References

- Abadie, Alberto (2003). “Semiparametric Instrumental Variable Estimation of Treatment Response Models”. In: *Journal of Econometrics* 113.2, pp. 231–263.
- Agarwal, Sumit, Shashwat Alok, Pulak Ghosh, and Sudip Gupta (2020). “Financial Inclusion and Alternate Credit Scoring for the Millennials: Role of Big Data and Machine Learning in Fintech”. Working Paper.
- Agrawal, Ajay, Joshua Gans, and Avi Goldfarb (2018). *Prediction Machines: The Simple Economics of Artificial Intelligence*. Boston, MA: Harvard Business Review Press.
- Bank for International Settlements (2019). *Big Tech in Finance: Opportunities and Risks*. BIS Annual Economic Report. Chap. III.
- Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace (2022). “Consumer-Lending Discrimination in the FinTech Era”. In: *Journal of Financial Economics* 143.1, pp. 30–56.
- Berg, Tobias, Valentin Burg, Ana Gombović, and Manju Puri (2020). “On the Rise of FinTechs: Credit Scoring Using Digital Footprints”. In: *The Review of Financial Studies* 33.7, pp. 2845–2897.
- Berg, Tobias, Valentin Burg, Jan Keil, and Manju Puri (2022). “On the Rise of Payment Firms”. Mimeo.
- Berg, Tobias, Andreas Fuster, and Manju Puri (2022). “FinTech Lending”. In: *Annual Review of Financial Economics* 14, pp. 187–207.
- Bhutta, Neil, Aurel Hizmo, and Daniel Ringo (2022). *How Much Does Racial Bias Affect Mortgage Lending? Evidence from Human and Algorithmic Credit Decisions*. FEDS Working Paper 2022-67.
- Blattner, Laura and Scott Nelson (2021). “How Costly is Noise? Data and Disparities in Consumer Credit”. arXiv preprint arXiv:2105.07554.
- Bonatti, Alessandro and Gonzalo Cisternas (2020). “Consumer Scores and Price Discrimination”. In: *The Review of Economic Studies* 87.2, pp. 750–791.

- Brunnermeier, Markus and Jonathan Payne (2022). “Platforms, Tokens, and Interoperability”. Unpublished Working Paper.
- Buchak, Greg, Gregor Matvos, Tomasz Piskorski, and Amit Seru (2018). “Fintech, Regulatory Arbitrage, and the Rise of Shadow Banks”. In: *Journal of Financial Economics* 130.3, pp. 453–483.
- Chen, Shuang, Clemens Sialm, and David X. Xu (2025). “The Growth and Performance of Artificial Intelligence in Asset Management”. Working Paper.
- Chen, Tao, Yi Huang, Chen Lin, and Zixia Sheng (2022). “Finance and Firm Volatility: Evidence from Small Business Lending in China”. In: *Management Science* 68.3, pp. 2226–2249.
- Chen, Xia, Qiang Cheng, Guanqun Dou, and Shuai Shao (2026). “The Value of Man in AI + Man: Field Evidence from Small Business Lending”. Working Paper.
- Choi, James J., Dong Huang, Zhishu Yang, and Qi Zhang (2025). *How Good is AI at Twisting Arms? Experiments in Debt Collection*. NBER Working Paper 33669.
- Cornelli, Giulio, Jon Frost, Leonardo Gambacorta, P. Raghavendra Rau, Robert Wardrop, and Tania Ziegler (2023). “Fintech and Big Tech Credit: Drivers of the Growth of Digital Lending”. In: *Journal of Banking & Finance* 148.
- Dai, Lili, Jianlei Han, Jing Shi, and Bohui Zhang (2022). “Digital Footprints as Collateral for Debt Collection”. Available at SSRN 4135159.
- Di Maggio, Marco, Dimuthu Ratnadiwakara, and Don Carmichael (2022). *Invisible Primes: Fintech Lending with Alternative Data*. NBER Working Paper w29840.
- Di Maggio, Marco and Vincent Yao (2021). “FinTech Borrowers: Lax Screening or Cream-Skimming?” In: *The Review of Financial Studies* 34.10, pp. 4565–4618.
- Dobbie, Will and Paige Marta Skiba (2013). “Information Asymmetries in Consumer Credit Markets: Evidence from Payday Lending”. In: *American Economic Journal: Applied Economics* 5.4, pp. 256–282.

- Foucault, Thierry, Leonardo Gambacorta, Wei Jiang, and Xavier Vives (2025). *Artificial Intelligence in Finance*. The Future of Banking. CEPR Press.
- Frost, Jon, Leonardo Gambacorta, Yi Huang, Hyun Song Shin, and Pablo Zbinden (2019). “BigTech and the Changing Structure of Financial Intermediation”. In: *Economic Policy* 34.100, pp. 761–799.
- Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther (2022). “Predictably Unequal? The Effects of Machine Learning on Credit Markets”. In: *The Journal of Finance* 77.1, pp. 5–47.
- Fuster, Andreas, Matthew Plosser, Philipp Schnabl, and James Vickery (2019). “The Role of Technology in Mortgage Lending”. In: *The Review of Financial Studies* 32.5, pp. 1854–1899.
- Gambacorta, Leonardo, Yiping Huang, Zhenhua Li, Han Qiu, and Shu Chen (2023). “Data versus Collateral”. In: *Review of Finance* 27.2, pp. 369–398.
- Gambacorta, Leonardo, Fabiana Sabatini, and Stefano Schiaffi (2025). *Artificial Intelligence and Relationship Lending*. BIS Working Paper 1244.
- Gertler, Paul, Brett Green, Renping Li, and David Sraer (2025). *The Welfare Benefits of Pay-As-You-Go Financing*. NBER Working Paper 33484.
- Ghosh, Pulak, Boris Vallee, and Yao Zeng (2025). “FinTech Lending and Cashless Payments”. In: *The Journal of Finance*.
- Gopal, Manasa and Philipp Schnabl (2022). “The Rise of Finance Companies and FinTech Lenders in Small Business Lending”. In: *The Review of Financial Studies* 35.11, pp. 4859–4901.
- Hau, Harald, Yi wHuang, Hongzhe Shan, and Zixia Sheng (2019). “How FinTech Enters China’s Credit Market”. In: *AEA Papers and Proceedings* 109, pp. 60–64.
- Howes, Cooper and Gregory Weitzner (2025). “Bank Information Production Over the Business Cycle”. In: *The Review of Economic Studies*.

- Huang, Yi, Sung Kwan Lee, and Bernard Yeung (2025). "Learning from AI-Finance: A Selected Synopsis of the AI-Fintech Conference in Aug 2024 in Shenzhen". In: *Journal of Finance and Data Science* 10.
- Huang, Yiping, Longmei Zhang, Zhenhua Li, Han Qiu, Tao Sun, and Xue Wang (2020). *Fintech Credit Risk Assessment for SMEs: Evidence from China*. IMF Working Paper 20/193.
- Jansen, Mark, Fabian Nagel, Constantine Yannelis, and Anthony Lee Zhang (2025). "Data and Welfare in Credit Markets". In: *Journal of Financial Economics* 174, p. 104171.
- Li, Jian and Stefano Pegoraro (2026). "Borrowing from a Bigtech Platform". In: *Review of Financial Studies*. forthcoming.
- Liu, Lei, Guangli Lu, and Wei Xiong (2022). "The Big Tech Lending Model". NBER Working Paper.
- Luohan Academy (2019). *Digital Technology and Inclusive Growth*.
- Stulz, René (2019). "FinTech, BigTech, and the Future of Banks". In: *Journal of Applied Corporate Finance* 31.4, pp. 86–97.
- Vasicek, Oldrich (1987). *Probability of Loss on Loan Portfolio*. Tech. rep. KMV Corporation.
- Yannelis, Constantine and Anthony Lee Zhang (2023). "Competition and Selection in Credit Markets". In: *Journal of Financial Economics* 150.2, p. 103710.

A Robustness for Section 3: Additional Specifications

This appendix reports four robustness exercises for the city-level regression in Section 3 (Table 2). The first three address identification concerns—whether the estimated relationship is driven by the Covid shock, by reverse causality, or by small-city outliers. The fourth replicates the analysis using a stricter 90-day delinquency definition.

A.1 Alternative samples, timing, and weights

Table A1 reports three exercises side by side, with two columns per exercise that correspond to specifications without and with GDP controls.

Excluding the Covid year (columns I–II). Restricting the sample to 2018–2019 drops 2020, the year in which the Covid shock is most likely to drive both the platform’s probability of default and real-economy defaults. With GDP controls (column II), the coefficient on defaulters is -0.034^* , statistically significant at the 10% level and comparable in magnitude to the baseline estimate of -0.027^{**} in Table 2. The no-GDP column (column I) is similar in sign and magnitude but not statistically significant, primarily reflecting the smaller sample. Dropping the shock year does not overturn the negative relationship.

Lagged defaulters (columns III–IV). Replacing the contemporaneous defaulter rate with its one-year lag yields an economically small and statistically insignificant coefficient in both columns, in contrast with the significant negative contemporaneous coefficient in the baseline. When both terms enter the regression jointly (with GDP controls and city fixed effects), the contemporaneous coefficient is -0.036^* and the lagged coefficient is effectively zero, confirming that the signal comes from contemporaneous co-movement rather than from a lagged feedback loop. Because court defaulter registration involves a substantial legal processing lag after an underlying default, this contemporaneous structure is more consistent with joint exposure to local economic conditions than with platform tightening causing subsequent court defaults.

Weighted by platform loan volume (columns V–VI). The baseline specification gives equal weight to all cities regardless of how much platform lending they receive, which could bias the estimate toward small cities with thin exposure. We therefore reweight each city-year observation by that city’s average weekly platform loan amount (*total_amt*) in 2019, drawn from our city-week dataset.

The reweighted estimates are -0.030^{***} (column V) and -0.021^{**} (column VI): the negative relationship strengthens in statistical significance, ruling out the concern that the baseline is driven by noise in low-exposure cities.

Table A1: Section 2 robustness: alternative samples, timing, and weights

Dependent variable: probability of default, platform loans (30 days, percent)						
	Excluding Covid year		Lagged defaulters		Weighted by 2019 loans	
	(I)	(II)	(III)	(IV)	(V)	(VI)
Defaulters (per thousand)	-0.022 (0.014)	-0.034^* (0.018)			-0.030^{***} (0.009)	-0.021^{**} (0.010)
Defaulters (per thousand), $t - 1$			-0.007 (0.014)	-0.009 (0.014)		
log GDP per capita		-0.416 (0.308)		-1.334^{***} (0.284)		-0.854^{***} (0.200)
GDP growth rate		-0.018 (0.023)		-0.005 (0.008)		-0.029^{***} (0.007)
Year 2018	-0.962^{***} (0.033)	-1.030^{***} (0.053)			-0.900^{***} (0.027)	-0.906^{***} (0.035)
Year 2020			0.605^{***} (0.032)	0.587^{***} (0.048)	0.736^{***} (0.024)	0.537^{***} (0.041)
City fixed effects	Y	Y	Y	Y	Y	Y
Adjusted R^2	0.899	0.892	0.915	0.939	0.921	0.922
N	538	360	538	376	813	592

Note: This table reports three robustness exercises for the city-level regression in Section 3 (Table 2). All specifications include city fixed effects; standard errors in parentheses. Columns (I)–(II) drop 2020 to check that the estimated relationship is not driven by the Covid shock. Columns (III)–(IV) replace the contemporaneous defaulter rate with its one-year lag; the lagged term is not statistically significant, indicating that the relationship is contemporaneous. When both the contemporaneous and lagged defaulter rates enter the regression jointly (with GDP controls and city fixed effects), the contemporaneous coefficient is -0.036^* while the lagged coefficient is indistinguishable from zero, confirming that the signal comes from contemporaneous co-movement. Columns (V)–(VI) weight each city-year observation by that city’s average weekly platform loan amount in 2019 (*total_amt*), so the estimate reflects cities with material platform exposure more heavily. $*** p < 0.01$, $** p < 0.05$, $* p < 0.1$.

A.2 Alternative PD definition: 90-day delinquency

Following Section 3, we apply the same analysis on 90 days probability of default of AI-based platform loans and report in Table A2. It shows the relationship between AI-based platform credit risk and real-economy risk is consistent with the analysis of 30 days probability of default rate in

Section 3.

Table A2: Fintech credit risk at the city level (90-day PD)

	Dep. var.: probability of default (90 days, %)				
	(I)	(II)	(III)	(IV)	(V)
Defaulters (per thousand)	0.020** (0.006)	-0.011** (0.004)	-0.018* (0.009)	-0.013** (0.005)	-0.021* (0.010)
GDP per capita				-0.125*** (0.037)	-0.716*** (0.196)
GDP growth rate				-0.010 (0.007)	-0.003 (0.007)
Year 2018		-0.781*** (0.038)	-0.786*** (0.026)	-0.781*** (0.047)	-0.862*** (0.037)
Year 2020		0.539*** (0.037)	0.541*** (0.020)	0.474*** (0.052)	0.522*** (0.038)
City fixed effects		Y	Y	Y	Y
Adj. R^2	0.011	0.601	0.840	0.567	0.838
N	818	818	818	655	655

Note: Dependent variable is 90-day probability of default (%). Columns (II)–(V) include city FEs; year dummies use 2019 as reference. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

B Robustness for Section 4: Level Specifications

Tables A4 and A5 present additional specifications for the dynamic credit adjustment analysis, using the level of LoC change (in RMB) instead of the IHS transformation. The qualitative patterns are consistent with the main results.

Inference with few time periods and the timing of the PD signal. The regressor of interest in Section 4, the previous month’s platform-wide PD, varies only across the months of the panel. Standard errors clustered at the borrower level therefore overstate precision. Table A3 re-estimates every column of Tables 6 to 9 on the sample of admitted borrower-months from May 2019 to December 2020, using the PD of the preceding calendar month, and reports standard errors clustered by borrower, clustered by month, two-way clustered, and Driscoll-Kraay. Three results survive inference at the month level: credit lines and interest rates are adjusted less often when aggregate risk is higher; conditional on an expansion, credit lines expand by more when risk is higher; and conditional on the borrower’s credit profile, higher risk is associated with a lower interest-rate change (column VI of Table 8) and with smaller unconditional rate increases and decreases (columns I, II, IV and V of Table 9). The net change in credit lines (IHS(Δ LoC), columns IV to VI of Table 6) and the contraction-margin coefficients are not statistically distinguishable from zero once the few-clusters problem is taken into account. A collapsed monthly time series (20 observations) with Newey-West standard errors gives the same picture.

Table A3: Dynamic credit adjustment: alternative standard errors for the coefficient on lagged PD

	(I)	(II)	(III)	(IV)	(V)	(VI)
<i>Table 4: Is LoC changed (I–III); IHS(ΔLoC) (IV–VI)</i>						
PD, previous month (ppt)	−0.145**	−0.154**	−0.158**	−0.197	−0.205	0.161
SE clustered by borrower	(0.001)	(0.001)	(0.001)	(0.006)	(0.007)	(0.008)
SE clustered by month	[0.054]	[0.056]	[0.058]	[0.536]	[0.569]	[0.630]
SE two-way borrower and month	{0.052}	{0.053}	{0.055}	{0.517}	{0.547}	{0.605}
SE Driscoll-Kraay	⟨0.051⟩	⟨0.055⟩	⟨0.057⟩	⟨0.405⟩	⟨0.401⟩	⟨0.497⟩
No. observations	3,535,758	3,064,334	3,064,334	3,535,758	3,064,334	3,064,334
<i>Table 5: IHS(ΔLoC) when LoC increases (I–III); decreases (IV–VI)</i>						
PD, previous month (ppt)	0.381**	0.380**	0.369**	−0.340*	−0.380*	−0.307
SE clustered by borrower	(0.003)	(0.004)	(0.004)	(0.006)	(0.007)	(0.006)
SE clustered by month	[0.159]	[0.165]	[0.154]	[0.192]	[0.186]	[0.194]
SE two-way borrower and month	{0.140}	{0.145}	{0.135}	{0.153}	{0.146}	{0.153}
SE Driscoll-Kraay	⟨0.137⟩	⟨0.141⟩	⟨0.131⟩	⟨0.194⟩	⟨0.190⟩	⟨0.201⟩
No. observations	705,305	632,302	632,302	271,183	224,464	224,464
<i>Table 6: Is rate changed (I–III); Δ rate, ppt (IV–VI)</i>						
PD, previous month (ppt)	0.015	0.014	0.027	0.125	0.128	−0.146**
SE clustered by borrower	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.002)
SE clustered by month	[0.064]	[0.067]	[0.060]	[0.116]	[0.122]	[0.053]
SE two-way borrower and month	{0.062}	{0.064}	{0.058}	{0.112}	{0.118}	{0.051}
SE Driscoll-Kraay	⟨0.042⟩	⟨0.044⟩	⟨0.035⟩	⟨0.100⟩	⟨0.106⟩	⟨0.060⟩
No. observations	3,535,758	3,064,334	3,064,334	3,535,758	3,064,334	3,064,334
<i>Table 7: Δ rate when rate increases (I–III); decreases (IV–VI)</i>						
PD, previous month (ppt)	−1.097***	−1.180***	−0.403	1.273***	1.332***	−0.147
SE clustered by borrower	(0.031)	(0.036)	(0.030)	(0.016)	(0.017)	(0.019)
SE clustered by month	[0.219]	[0.212]	[0.277]	[0.245]	[0.243]	[0.443]
SE two-way borrower and month	{0.165}	{0.160}	{0.207}	{0.181}	{0.179}	{0.326}
SE Driscoll-Kraay	⟨0.133⟩	⟨0.137⟩	⟨0.282⟩	⟨0.188⟩	⟨0.194⟩	⟨0.408⟩
No. observations	40,469	30,606	30,606	117,584	99,771	99,771
Shop controls	N	Y	Y	N	Y	Y
Lagged log LoC and rate	N	N	Y	N	N	Y
Individual FE	Y	Y	Y	Y	Y	Y

Note: Each panel re-estimates the corresponding table in the main text and reports the coefficient on the previous month's platform-wide PD (in percentage points) with four alternative standard errors: clustered by borrower (parentheses, as in the main text), clustered by month (brackets, 20 clusters), two-way clustered by borrower and month (braces), and Driscoll-Kraay with two lags (angle brackets). Because PD varies only across the 20 months of the regression sample, the month-clustered and Driscoll-Kraay standard errors are the relevant ones for inference. Stars are based on the month-clustered standard errors with $t(20)$ critical values: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A4: Credit adjustment for individual borrowers relative to past PD

	Is LoC changed (0/1)			Δ LoC (RMB)		
	(I)	(II)	(III)	(IV)	(V)	(VI)
PD, prev. month (ppt.)	-0.144*** (0.001)	-0.144*** (0.001)	-0.143*** (0.001)	-252*** (12.4)	-156*** (13.8)	261*** (15.3)
Log LoC, prev. month			-0.015*** (0.000)			-2,009*** (17.0)
Interest, prev. month (ppt.)			-0.002*** (0.000)			-15*** (2.8)
Shop controls	N	Y	Y	N	Y	Y
Individual FE	Y	Y	Y	Y	Y	Y
R^2	0.155	0.162	0.163	0.041	0.049	0.072
N	3,382,776	2,932,390	2,932,390	3,382,776	2,932,390	2,932,390

Note: Baseline regressions of LoC adjustment on lagged PD. Shop controls include log sales value and log sales orders. Clustered SEs (borrower level) in parentheses. *** $p < 0.001$.

Table A5: Credit adjustment for individual borrowers is asymmetric

Dependent variable: Sample	Change in LoC (level, RMB)					
	When LoC increases (Δ LoC > 0)			When LoC decreases (Δ LoC < 0)		
	(I)	(II)	(III)	(IV)	(V)	(VI)
PD, prev. month avg. (ppt.)	4,300.0*** (55.4)	4,435.2*** (60.5)	4,404.6*** (65.1)	-1,299.2*** (99.7)	-1,729.7*** (112.4)	-1,342.3*** (105.7)
Log LoC, prev. month			-716.5*** (27.8)			-11,600*** (125.7)
Interest, prev. month (ppt.)			-201.3*** (9.0)			-396.4*** (13.6)
Shop controls	N	Y	Y	N	Y	Y
Individual FE	Y	Y	Y	Y	Y	Y
R^2	0.501	0.512	0.513	0.652	0.674	0.730
No. observations	691,339	625,065	625,065	326,398	282,082	282,082

Note: This table presents the regression results examining the relationship between aggregate risk and adjustment in line of credit (LoC). Standard errors, reported in parentheses, are clustered at the individual borrower level. Shop controls include the log sales value (RMB) and log sales orders. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

C Robustness for Section 5: Instrumental Variable Analysis

This appendix collects additional analyses for the instrumental-variable framework in Section 5. Subsection C.1 reports robustness using alternative measures and instruments. Subsection C.2 reports the validity tests (reduced-form progressive controls, lead placebo) and identification diagnostics (cohort event study, Wald-IV ratio, complier profile). Subsection C.3 repeats the analysis at the loan level instead of the borrower-month level. Subsection C.4 reports heterogeneity by shop tenure and the owner’s digital age that complements the main-text mechanism analysis.

C.1 Alternative Measures and Alternative Instruments

Alternative measure: number of buyers. Table A6 substitutes $\log(\text{number of buyers})$ for $\log(\text{sales})$ as the measure of platform dependency. The instrument set, controls, and fixed effects match the main-text Table 14. The 2SLS coefficient on $\log(\text{buyers})$ is consistently negative and strongly significant across the one-, two-, and three-instrument specifications, mirroring the main-text result that higher platform activity reduces default. The OLS estimate is again attenuated relative to 2SLS, consistent with the LATE heterogeneity story.

Alternative instrument: Consumer Protection Bail. Table A7 replaces the Gold Seller instrument with the log of the seller’s Consumer Protection Bail deposit, both as a current-month value and as a three-month average. Both deposit IVs deliver strong first-stage F-statistics and produce 2SLS estimates of platform dependency on default that are larger in magnitude than the Gold Seller IV, suggesting that Gold Seller may be a conservative instrument.

Long-period sales. As a further robustness check, Table A8 reports results using 3-month cumulative sales instead of monthly sales. The qualitative pattern is preserved across the longer measurement window.

Table A6: Platform dependency reduces default: Alternative measure

	Default dummy ($\times 100$)			Default amount share (%)		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Two-Stage Least Squares</i>						
Log (number of buyers)	−0.087*** (0.004)	−0.082*** (0.006)	−0.096*** (0.010)	−0.067*** (0.003)	−0.062*** (0.005)	−0.075*** (0.008)
Log (loan amount)	0.004 (0.003)	−0.004 (0.004)	−0.006 (0.007)	−0.000 (0.002)	−0.008* (0.004)	−0.010 (0.006)
Number of loans	0.016*** (0.001)	0.015*** (0.002)	0.014*** (0.002)	0.009*** (0.001)	0.009*** (0.001)	0.009*** (0.002)
Shop open years	−0.007*** (0.002)	−0.006* (0.003)	−0.008+ (0.004)	−0.005*** (0.002)	−0.005* (0.002)	−0.007+ (0.004)
<i>Panel B: First Stage for Log (number of buyers)</i>						
Gold seller dummy (0/1)	1.674*** (0.005)	1.235*** (0.005)	1.103*** (0.008)	1.674*** (0.005)	1.235*** (0.005)	1.103*** (0.008)
Gold seller (lag 1)		0.837*** (0.005)	0.669*** (0.007)		0.837*** (0.005)	0.669*** (0.007)
Gold seller (lag 2)			0.473*** (0.008)			0.473*** (0.008)
F-Statistic	309,231	84,908	26,900	309,231	84,908	26,900
Adjust R^2	0.329	0.353	0.368	0.329	0.353	0.368
<i>Panel C: Ordinary Least Squares</i>						
Log (number of buyers)	−0.052*** (0.002)	−0.053*** (0.003)	−0.065*** (0.005)	−0.043*** (0.002)	−0.044*** (0.003)	−0.052*** (0.005)
Gold seller dummy (0/1)	−0.058*** (0.007)	−0.064*** (0.012)	−0.073*** (0.019)	−0.040*** (0.006)	−0.052*** (0.010)	−0.057*** (0.017)
Gold seller (lag 1)		0.009 (0.013)	−0.023 (0.019)		0.018 (0.011)	−0.014 (0.017)
Gold seller (lag 2)			0.036+ (0.021)			0.028 (0.018)
Loan Controls	YES	YES	YES	YES	YES	YES
Shop and Borrower Controls	YES	YES	YES	YES	YES	YES
Industry FE	YES	YES	YES	YES	YES	YES
City-Month FE	YES	YES	YES	YES	YES	YES
Observations	2,266,337	1,025,290	462,275	2,266,337	1,025,290	462,275

Note: Replaces log(sales) with log(number of buyers) as the platform-dependence measure. Panels and columns follow Table 14. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, + $p < 0.10$.

Table A7: Platform dependency reduces default: Alternative IV

	Default dummy ($\times 100$)		Default amount share (%)	
	(1)	(2)	(3)	(4)
<i>Panel A: Two-Stage Least Squares</i>				
Log (sales)	-0.320*** (0.022)	-0.300*** (0.023)	-0.268*** (0.018)	-0.250*** (0.018)
<i>Panel B: First Stage for Log (sales)</i>				
Log (deposit amount)	0.051*** (0.001)		0.051*** (0.001)	
Log (3-month average deposit amount)		0.051*** (0.001)		0.051*** (0.001)
F-Statistic	21,093	20,095	21,093	20,095
Adjust R^2	0.270	0.269	0.270	0.269
<i>Panel C: Ordinary Least Squares</i>				
Log (sales)	-0.059*** (0.002)	-0.059*** (0.002)	-0.048*** (0.002)	-0.048*** (0.002)
Log (deposit amount)	-0.013*** (0.001)		-0.011*** (0.001)	
Log (3-month average deposit amount)		-0.012*** (0.001)		-0.010*** (0.001)
Loan Controls	YES	YES	YES	YES
Shop and Borrower Controls	YES	YES	YES	YES
Industry FE	YES	YES	YES	YES
City-Month FE	YES	YES	YES	YES
Observations	2,266,337	2,266,337	2,266,337	2,266,337

Note: Uses the log of the Consumer Protection Bail deposit as an alternative instrument for log(sales), either current-month or 3-month average. Loan controls exclude log(deposit) when it enters as instrument. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, + $p < 0.10$.

Table A8: Platform dependency reduces default: Long period measure

	Default dummy ($\times 100$)			Default amount share (%)		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Two-Stage Least Squares</i>						
Log (3-months sales)	-0.345*** (0.015)	-0.312*** (0.016)	-0.326*** (0.024)	-0.149*** (0.007)	-0.132*** (0.009)	-0.146*** (0.014)
<i>Panel B: First Stage for Log (Sales)</i>						
Gold seller dummy (0/1)	2.419*** (0.073)	1.507*** (0.034)	1.157*** (0.032)	2.419*** (0.073)	1.507*** (0.034)	1.157*** (0.032)
Gold seller (lag 1)		1.543*** (0.048)	1.055*** (0.027)		1.543*** (0.048)	1.055*** (0.027)
Gold seller (lag 2)			1.131*** (0.032)			1.131*** (0.032)
F-Statistic	135,811	42,313	14,972	135,811	42,313	14,972
Adjust R^2	0.398	0.424	0.435	0.398	0.424	0.435
Loan Controls	YES	YES	YES	YES	YES	YES
Shop and Borrower Controls	YES	YES	YES	YES	YES	YES
Industry FE	YES	YES	YES	YES	YES	YES
City-Month FE	YES	YES	YES	YES	YES	YES
Observations	640,303	321,691	161,051	640,303	321,691	161,051

Note: This table shows the 2SLS regressions for default variables using 3-month sales as the measure of platform dependency. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

C.2 IV Validity and Identification Diagnostics

Reduced-form regressions with progressive controls. Table A9 reports reduced-form regressions of default on the Gold Seller indicator, progressively adding controls from a fixed-effects-only specification (column 1) through to a full set including log(sales), log(buyers), and log(orders) (column 6). The Gold Seller coefficient is large and significant in columns 1–3 but shrinks substantially once log(sales) is included in column 4 – consistent with the platform-dependency channel operating primarily through observable sales rather than through an unobserved direct channel. Adding log(buyers) and log(orders) in columns 5–6 yields essentially no further attenuation, indicating that contemporaneous platform-activity measures are close substitutes.

Table A9: IV validity: reduced-form with progressive controls

	(1) FE only	(2) + Loan	(3) + Full ctrl	(4) + log(sales)	(5) + log(buyers)	(6) + log(orders)
Gold seller dummy	−0.172*** (0.006)	−0.157*** (0.007)	−0.145*** (0.007)	−0.023** (0.008)	−0.020* (0.008)	−0.021** (0.008)
Log (sales)				−0.057*** (0.002)	−0.042*** (0.003)	−0.041*** (0.003)
Log (buyers)					−0.022*** (0.003)	−0.017*** (0.003)
Log (orders)						−0.012*** (0.004)
Observations	2,266,337	2,266,337	2,266,337	2,266,337	2,266,337	2,266,337

Note: Reduced-form regressions of default on Gold Seller status with progressively added controls. Log(sales) absorbs most of the Gold Seller–default association; adding log(buyers) and log(orders) does not provide meaningful further reduction, indicating that contemporaneous platform-activity measures are close substitutes. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, + $p < 0.10$.

Lead placebo test. Table A10 reports a placebo specification in which the lead (one-month-ahead) Gold Seller dummy is added to the reduced-form regression. The lead enters strongly when no other Gold Seller terms are included (column 1), reflecting persistence of Gold Seller status across months. Once the current Gold Seller (column 2) and a one-month lag (column 3) are added, the lead is partially attenuated; further conditioning on log(sales) (column 4) reduces the lead by an additional ~24%. The pattern is consistent with a contemporaneous causal channel through sales, with the residual lead capturing the seller-improvement trajectory documented in

the event-study figures of the main text.


Table A10: Lead Gold Seller placebo

	(1) Lead only	(2) +Current	(3) +Current+Lag	(4) +Sales	(5) +All activity
Gold seller (lead 1)	−0.097*** (0.006)	−0.078*** (0.007)	−0.084*** (0.011)	−0.059*** (0.007)	−0.059*** (0.007)
Gold seller (current)		−0.035*** (0.008)	−0.036*** (0.011)	−0.009 (0.008)	−0.010 (0.008)
Gold seller (lag 1)			0.002 (0.013)		
Log (sales)				−0.018*** (0.002)	−0.016*** (0.003)
Log (buyers)					0.001 (0.003)
Log (orders)					−0.009** (0.003)
Observations	980,936	980,936	442,026	980,936	980,936

Note: Columns progressively control for current Gold Seller status and contemporaneous platform activity. The lead coefficient partially attenuates (cols 2–3) but remains significant; adding log(sales) in col 4 reduces the lead by a further 24%, and additional activity measures (cols 5) provide no marginal reduction. The residual lead coefficient reflects the seller-improvement trajectory in the months preceding first Gold Seller certification — the same dynamic visible in the event study (Figure A1). *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, + $p < 0.10$.

Cohort-specific event study. We split the treated population by the calendar year of first Gold Seller certification (2018, 2019, and 2020) and re-estimate the event study within each cohort using a common pool of never-treated controls. Figure A1 reports the Sun–Abraham (2021) interaction-weighted estimates. The first-stage panel shows that all three cohorts produce a clean V-shape with a sharp positive jump in log(sales) at $k = 0$ that decays over the following two to three months. The reduced-form panel for default is more heterogeneous: the 2019 cohort shows the cleanest negative response at $k = 0$ (consistent with the pooled main-text estimate), whereas the 2018 and 2020 cohorts show smaller and noisier responses. The pooled estimate in the main text is therefore an average across cohort-level effects of varying magnitude, but the qualitative direction (negative at $k = 0$) is preserved.

Figure A1: Cohort-specific event study around the first Gold Seller certification




Note: Treated borrowers are split by the calendar year of first Gold Seller certification (2018, 2019, 2020); the never-treated control group is common across cohorts. Estimator is Sun–Abraham (2021) interaction-weighted, with borrower and city×month fixed effects. Bars show 95% confidence intervals based on standard errors clustered at the borrower level. Baseline period $k = -1$ is omitted by construction.

Wald-IV ratio by horizon. For each event time k we compute the horizon-specific Wald-IV estimate $\hat{\beta}_k^{\text{RF}} / \hat{\beta}_k^{\text{FS}}$. Under the exclusion restriction the ratio should be approximately constant across k and equal to the structural parameter. Figure A2 shows the ratio decomposed by event time, with the Section 5 IV reference of -0.075 overlaid. The pre-event period ($k = -6$) gives a ratio close to -0.34 , while $k = 0$ gives a ratio near -0.07 matching the main-text estimate; the ratio becomes unstable for $k \geq 3$ where the first stage has decayed. This decomposition shows that the main-text IV estimate is most representative of the contemporaneous-shock effect rather than longer-horizon co-movements.

Complier profile (Abadie 2003). Table A11 reports Abadie (2003) κ -weighted means of borrower characteristics for the complier subpopulation – those whose log(sales) is induced to cross the median by Gold Seller status. The complier share is $\hat{\pi}_c \approx 0.54$. Compliers have lower sales (by 10%) and fewer buyers (by 16%) than the full sample, but a 24% higher baseline default rate. This indicates that the IV identifies the local effect for marginal, platform-dependent, higher-risk sellers – exactly the subgroup for which the platform-incentive channel emphasised in the main

Figure A2: Wald-IV ratio by event time



Note: For each event time k , the figure plots $\hat{\beta}_k^{\text{RF}} / \hat{\beta}_k^{\text{FS}}$, the horizon-specific Wald-IV estimate of the effect of $\log(\text{sales})$ on default (in percentage points per log-point sales). Under the exclusion restriction, the ratio should be approximately constant across k and equal to the structural parameter. The dotted blue line marks the main-text IV reference of -0.075 from Table 14. Periods with $|\hat{\beta}_k^{\text{FS}}| < 0.05$ are dropped because the ratio is unstable when the first stage has decayed. Bars show 95% confidence intervals.

text would be expected to bind most strongly.

Table A11: Complier profile: characteristics of marginal sellers identified by the Gold Seller IV

Variable	Full sample	Complier	Ratio	Difference
Log (sales)	9.450	8.500	0.900	-0.950
Log (number of buyers)	4.317	3.649	0.845	-0.668
Log (loan amount)	10.403	10.168	0.977	-0.235
Number of active loans	3.845	3.631	0.944	-0.214
Shop open years	6.097	6.122	1.004	+0.025
Log (online time)	10.107	9.919	0.981	-0.188
Shop rating	0.991	0.990	0.999	-0.001
Forge rate	14.095	14.322	1.016	+0.227
Years used digital payment	7.418	7.441	1.003	+0.023
Is car owner	0.607	0.592	0.975	-0.015
Is property owner	0.932	0.929	0.997	-0.003
Log (deposit)	4.692	4.351	0.927	-0.341
P4P participant	0.839	0.813	0.969	-0.026
Default (pp)	0.205	0.254	1.239	+0.049
Complier share $\hat{\pi}_c$	0.544			
Sample (after PS trim)	2,113,300 borrower-months			

Note: Abadie (2003) κ -weighted means of borrower characteristics for the complier subpopulation, defined as borrower-months induced to cross the median log(sales) by Gold Seller status. The propensity score $P(Z = 1 \mid W)$ is estimated by logit on shop tenure, online time, demographics, deposit, forge rate, good rate, and industry FE; observations with $\hat{p} \notin (0.01, 0.99)$ are trimmed. Compliers have lower sales (by 10%) and buyers (by 16%) than the full sample, but a 24% higher baseline default rate, indicating that the IV identifies a marginal, platform-dependent, higher-risk subgroup.

C.3 Robustness at the Loan Level

The main analysis in Section 5 uses a borrower-month panel: each observation is one calendar month during which a borrower has at least one active loan. As a robustness check on the unit of observation, this subsection re-estimates the analysis at the *loan level*: each observation is one completed loan, the dependent variable is whether the loan ultimately defaulted (0/100), the explanatory variable of interest is the average $\log(\text{sales})$ over the loan's active life, and the instrument is the fraction of loan-life months during which the borrower was a Gold Seller. Loan-level analysis avoids potential confounding from time-aggregation choices and delivers the loan-level interpretation of the platform-dependency effect.

Table A12 reports descriptive statistics for the loan-level analysis sample. The sample contains 2,418,347 completed loans on 367,036 borrowers, with an overall default rate of 2.2%, a median loan life of seven months, and an annualized interest rate averaging 15.0%. Shop-level variables are averaged over each loan's active months; Gold Seller exposure is captured both as the fraction of loan-life months with certification (mean 0.18) and as a binary indicator equal to 1 if the borrower was ever a Gold Seller during the loan's life (mean 0.29).

Across all five regression tables in this subsection, the qualitative pattern from the borrower-month analysis is preserved: higher platform sales reduce the probability of loan default, the effect is concentrated among platform-dependent sellers, and the IV consistently delivers a larger absolute effect than OLS.

Loan-level main IV. Table A13 reports the loan-level main-IV result. A one log-point increase in life-averaged sales reduces the loan's default probability by 0.47 percentage points; the first-stage F-statistic of 7.5×10^5 indicates an extremely strong instrument. The 2SLS estimate (-0.470) is roughly twice the OLS estimate (-0.260), an OLS-IV wedge in the same direction as the borrower-month evidence in Section 5.

Alternative measure: number of buyers. Table A14 substitutes life-averaged $\log(\text{buyers})$ for $\log(\text{sales})$ as the platform-dependence measure. The IV coefficient is -0.598 (SE 0.048), preserving the negative direction of the platform-dependence effect; the somewhat larger magnitude reflects that buyer counts move less than sales per unit of platform exposure.

Table A12: Descriptive statistics at the loan level

Variable	Mean	P25	Median	P75	Std. Dev.	N
<i>Loan information</i>						
Default (0/1)	0.022	0.000	0.000	0.000	0.147	2,418,347
Loan amount (yuan)	20,478	2,000	7,050	20,000	46,095	2,418,347
Loan life (months)	7.67	4.00	7.00	13.0	4.53	2,418,347
Interest rate (annualized, %)	14.95	13.69	14.96	15.82	2.24	2,418,347
Forge rate (credit score)	14.8	13.9	14.8	15.6	2.08	2,418,347
<i>Shop information (life-averaged)</i>						
Log (avg sales over loan life)	9.41	8.08	9.48	10.80	2.02	2,418,347
Log (avg buyers over loan life)	4.25	2.84	4.18	5.56	1.88	2,393,033
Frac (Gold Seller) over loan life	0.177	0.000	0.000	0.200	0.328	2,418,347
Any Gold Seller during loan life	0.287	0.000	0.000	1.000	0.452	2,418,347
Avg shop rating over loan life	0.99	0.99	0.99	1.00	0.02	2,418,347
Log (avg CS online time)	9.75	9.41	10.16	10.63	1.71	2,299,700
Avg shop open years	6.40	4.11	6.04	8.47	3.03	2,418,347
Frac P4P over loan life	0.802	1.000	1.000	1.000	0.394	2,418,347
<i>Borrower Characteristics</i>						
Age	35.7	29.0	33.0	39.0	9.79	366,967
Years used digital payment	7.20	5.00	7.00	9.00	3.04	367,036
Is property owner	0.909	1.000	1.000	1.000	0.288	366,278
Is car owner	0.587	0.000	1.000	1.000	0.492	366,278

Note: Descriptive statistics for the loan-level analysis sample used in Appendix C.3. Each observation is one completed loan. Shop-level variables (sales, buyers, Gold Seller status, shop rating, customer-service online time, shop open years, P4P participation) are averaged across the loan's active months from disbursement to end-of-life. Frac (Gold Seller) is the fraction of loan-life months during which the borrower held Gold Seller status; Any Gold Seller is a binary indicator equal to 1 if the borrower was Gold Seller in at least one month during the loan's life. The interest rate is the annualized contracted rate. Borrower-level variables (Age, digital payment years, property/car ownership) are reported on the unique-borrower basis.

Alternative instrument: ever Gold Seller during loan life. Table A15 replaces the continuous instrument (fraction of loan-life months as Gold Seller) with a binary indicator equal to 1 if the borrower was ever a Gold Seller during the loan's active life. The 2SLS estimate of -0.334 is modestly smaller in absolute magnitude than the main specification, as expected for the coarser instrument; the sign and significance are unchanged.

IV validity: reduced form with progressive controls. Table A16 reports the loan-level reduced-form of default on the Gold Seller life-fraction with progressively richer controls. The

Table A13: Platform dependency reduces default at the loan level

	First stage log(sales) ^{life}	2SLS default (pp)	OLS default (pp)
Log (sales) ^{life} (instrumented)		−0.470*** (0.037)	−0.260*** (0.030)
Frac (Gold Seller) ^{life}	2.729*** (0.012)		−0.574*** (0.130)
Loan + Borrower controls	YES	YES	YES
Industry + City×Disburse-month FE	YES	YES	YES
First-stage F	751,284	—	—
Adjust R ²	0.486	—	—
Observations	2,296,035	2,296,035	2,296,035

Note: Loan-level cross-section. Each observation is one loan with completed status (settled, defaulted, or written off). `frac_gold` is the fraction of months in the loan’s active life during which the borrower held Gold Seller status; `avg_log_sales` is the mean of log monthly sales over the loan’s life. Standard errors clustered at the borrower level. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, + $p < 0.10$.

Table A14: Loan-level alternative measure of platform dependency: number of buyers

	First stage log(buyers) ^{life}	2SLS default (pp)	OLS default (pp)
Log (buyers) ^{life}		−0.598*** (0.048)	−0.188*** (0.030)
Frac (Gold Seller) ^{life}	2.134*** (0.015)		
Loan + Borrower controls	YES	YES	YES
Industry + City×Disburse-month FE	YES	YES	YES
First-stage F	445,614	—	—
Observations	2,280,929	2,280,929	2,280,929

Note: Loan-level cross-section. Replaces life-averaged log(sales) with life-averaged log(number of buyers) as the platform-dependence measure. The instrument is the fraction of months in the loan’s active life during which the borrower held Gold Seller status. Loan, borrower, and shop controls match Table A13. Standard errors clustered at the borrower level. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, + $p < 0.10$.

Frac (Gold Seller) coefficient shrinks from −1.459 (FE only, column 1) to −0.574 once life-averaged log(sales) is added (column 4) – a 60% reduction that mirrors the borrower-month finding (Table A9) that the bulk of the Gold Seller–default association is mediated by observable platform activity. Adding log(buyers) in column 5 produces no further reduction.

Table A15: Loan-level alternative IV: ever Gold Seller during loan life

	First stage log(sales) ^{life}	2SLS default (pp)	OLS default (pp)
Log (sales) ^{life}		−0.334*** (0.039)	−0.304*** (0.030)
Any Gold Seller ^{life}	2.011*** (0.008)		
Loan + Borrower controls	YES	YES	YES
Industry + City×Disburse-month FE	YES	YES	YES
First-stage F	781,309	—	—
Observations	2,296,035	2,296,035	2,296,035

Note: Loan-level cross-section. Replaces the continuous instrument `frac_gold` (the fraction of loan-life months as Gold Seller) with a binary indicator equal to 1 if the borrower was ever a Gold Seller during the loan’s active life. Same outcome, controls, and fixed effects as Table A13. Standard errors clustered at the borrower level. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, + $p < 0.10$.

Table A16: Loan-level IV validity: reduced-form with progressive controls

	(1) FE only	(2) + Loan	(3) + Full ctrl	(4) + Sales	(5) + Buyers
Frac (Gold Seller) ^{life}	−1.459*** (0.096)	−1.268*** (0.099)	−1.282*** (0.101)	−0.574*** (0.130)	−0.560*** (0.131)
Log (sales) ^{life}				−0.260*** (0.030)	−0.242*** (0.034)
Log (buyers) ^{life}					−0.029 (0.034)
Industry + City×Disburse-month FE	YES	YES	YES	YES	YES
Observations	2,418,005	2,418,005	2,296,035	2,296,035	2,280,929

Note: Loan-level reduced-form regressions of default on the Gold Seller life-fraction, with progressively richer controls. Column (1) includes only fixed effects; col. (2) adds loan-level controls (loan amount, interest rate, forge rate, life months); col. (3) adds the full set of borrower and shop-life controls; col. (4) adds life-averaged log(sales); col. (5) adds life-averaged log(buyers). The Frac (Gold Seller) coefficient drops from −1.46 (FE only) to −0.57 when sales is added, mirroring the borrower-month finding (Table A9) that the bulk of the Gold Seller–default association is mediated by observable platform activity. Standard errors clustered at the borrower level. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, + $p < 0.10$.

Loan-level heterogeneity. Table A17 reports loan-level heterogeneity by three split criteria: life duration, P4P advertising participation, and loan size. The platform-dependence effect is five times stronger for loans of life > 2 months (−0.530) than for short loans (−0.107), consistent

with the intuition that longer loan exposure leaves more scope for the platform-incentive channel to operate. The effect is also slightly stronger for P4P participants and for smaller loans, both consistent with the mechanism evidence in Section 5.

Table A17: Loan-level heterogeneity: 2SLS estimate of $\log(\text{sales})^{\text{life}}$ on default

Subsample	Coefficient	First-stage F	N
Life $\leq$ 2 months	−0.107*** (0.013)	112,967	457,837
Life $>$ 2 months	−0.530*** (0.047)	664,071	1,837,016
P4P high	−0.512*** (0.043)	600,573	1,853,935
P4P low	−0.439*** (0.055)	114,208	441,058
Loan amount large	−0.430*** (0.032)	417,417	1,164,157
Loan amount small	−0.554*** (0.061)	299,056	1,130,972

Note: 2SLS estimates of life-averaged $\log(\text{sales})$ on the loan-default dummy multiplied by 100. Subsamples split as indicated. $\text{Frac}(\text{Gold Seller})^{\text{life}}$ is the instrument throughout. Standard errors clustered at the borrower level.

C.4 Heterogeneity by Shop Tenure and Owner's Digital Age

The main-text Table 16 reports heterogeneity in the platform-dependency effect by P4P participation, seller star rating, and an outstanding-balance cash-flow control. As a complement to those tenure-based proxies for platform-specific investment, Table A18 reports the analogous splits by shop tenure (years since the seller opened the shop) and the owner's digital age (years since the owner first registered as a consumer on the platform, prior to opening a shop). The estimated 2SLS effects are similar across short- and long-tenure subgroups, consistent with the interpretation that these variables capture general business or platform familiarity rather than platform-specific investment, and therefore provide a weaker discriminating signal than the P4P and star-rating proxies in the main text.

Table A18: Heterogeneity by shop tenure and owner's digital age

	Dep. var.: Default dummy ($\times 100$)	
	(1)	(2)
<i>Shop tenure</i>	Short	Long
Log (sales)	-0.069*** (0.007) [63, 183]	-0.063*** (0.007) [69, 159]
<i>N</i>	505,224	518,442
<i>Owner's digital age</i>	Short	Long
Log (sales)	-0.081*** (0.008) [50, 732]	-0.056*** (0.006) [81, 904]
<i>N</i>	413,882	609,781
Controls & FEs	Yes	

Note: 2SLS estimates of log(sales) on default ($\times 100$), splitting the sample at the median of the indicated variable. Same specification, controls, instruments and fixed effects as the main heterogeneity analysis (Table 16). Differences across tenure and digital-age subgroups are economically modest, consistent with these proxies capturing general business experience rather than platform-specific investment. First-stage *F*-statistics in brackets; standard errors in parentheses, clustered at the borrower level. *** $p < 0.001$.