

Fintech Access and Consumption Smoothing^{*}

Long Chen

Luohan Academy

Xavier Giroud

Columbia University, NBER, CEPR

Yingju Ma

University of Cambridge

Neng Wang

Cheung Kong Graduate School of Business

August 10, 2026

Abstract

Fintech credit through digital payment platforms reaches 500 million Chinese consumers, many of whom hold no traditional bank credit line. We use de-identified administrative data from Alipay covering 60 months of full payment activity, both online and in-person, across payment instruments and consumption categories, for a sample drawn from a near-billion-consumer population. This provides the first integrated view of consumer payment and credit behavior in China. Using the January 2020 COVID-19 outbreak as a quasi-exogenous liquidity shock and a matched difference-in-differences design on 434,398 consumers, we find that fintech credit users sustain 8.6% higher consumption over the following two years. The response is concentrated among consumers for whom an additional liquidity backstop should have high marginal value: low-spending consumers, consumers in less developed regions, consumers with limited active use of substitute credit, and recent entrants to the digital-payment ecosystem. The evidence points to a precautionary liquidity-buffer channel: an immediately usable credit line supports consumption smoothing even though 86% of users never carry an installment balance over the two post-shock years.

^{*}Long Chen is with Luohan Academy. Xavier Giroud is with Columbia University, NBER, and CEPR. Yingju Ma is with University of Cambridge. Neng Wang is with Cheung Kong Graduate School of Business. We benefited from the comments of Daron Acemoglu, Martin Beraja, Patrick Bolton, Christopher J. Palmer, Jonathan A. Parker, Nagpurnanand Prabhala, Yessenia Tellez, Robert M. Townsend, Constantine Yannelis, and participants at NASM 2024 and MFA 2026. We thank Luohan Academy, supported by Ant Group and Alibaba Group, which have helped create de-identified data that are selectively available to be used for academic research. We are grateful to the data scientists team at Luohan Academy for their assistance in constructing the datasets. Zhiren Hu and Boyang Zhang have provided outstanding research assistance at the early stage. The opinions expressed are those of the authors alone and do not represent Ant Group's or Alibaba Group's views. For comments and suggestions, please email ym459@cam.ac.uk. (First posted version: October 30, 2024. Click [here](#) for the latest version.)

Much of the past decade’s expansion of consumer credit has occurred outside the banking system, through credit lines embedded in digital payment apps. By mid-2020, one such product, Huabei, reached 500 million consumers in China, most of whom hold no bank credit card.¹ Credit on this scale, extended to consumers whom the formal financial system barely observes, confronts regulators with two questions: does payment-embedded credit help households sustain consumption, and does it lead them into debt they cannot manage?² This paper provides individual-level evidence on both questions, at the scale at which this credit actually operates.

The contract we study shows why the answers are not straightforward extensions of what is known about credit cards. Huabei is embedded in Alipay, one of the two dominant digital payment platforms in China: at checkout, the consumer chooses between the account balance, a linked bank card, and the credit line, so drawing on the line involves no application, delay, or transaction cost. Repayment within the monthly billing cycle is interest-free, the full balance is deducted automatically on the bill date unless the user actively elects installments, and only balances rolled into installments accrue interest, at an annual rate of roughly 15%. Credit limits are individually assigned and can be as low as 20 yuan; the average outstanding balance was about 2,000 yuan in 2020. Holding the line is free, and settling in full each cycle is costless, so access and borrowing are separated by the product’s design: a consumer can hold the line for years without ever paying interest.

We measure what this access does to consumption using de-identified administrative records from Alipay: monthly spending for each consumer from 2017 to 2021, online and in person, across payment instruments and merchant categories, together with credit-line usage, repayment, delinquency, and demographic records. The sample is restricted to consumers with high pre-shock platform activity, so that recorded payments capture the bulk of their spending. The design compares consumers who had activated Huabei before January 2020 with one-to-one matched non-users around the COVID-19 outbreak, a sudden economy-wide liquidity shock, in a matched difference-in-differences framework covering 434,398 consumers; matching is on two years of pre-shock consumption, age, gender, region, and the year of entry into digital payments, and estimation is by Poisson pseudo-maximum likelihood with individual and city-by-month fixed effects.

Access to the credit line supported consumption when the shock hit, and the support did not

¹By June 2020, Huabei served about 500 million users with an average outstanding balance of 2,000 yuan, making it the largest consumer credit service in China by both credit balance and number of users. Source: Ant Group’s IPO prospectus, Hong Kong Stock Exchange, October 26, 2020. Section 1.1 provides institutional detail.

²In the United States, the Consumer Financial Protection Bureau opened a market-monitoring inquiry into buy-now-pay-later products in December 2021 and published its findings in September 2022 (*Buy Now, Pay Later: Market Trends and Consumer Impacts*). In the United Kingdom, the Woolard Review (Financial Conduct Authority, 2021) recommended bringing interest-free short-term credit within the regulatory perimeter. In the European Union, the revised Consumer Credit Directive (Directive (EU) 2023/2225) extends consumer-credit protections to such products.

fade. Consumers with access sustain 9.0% higher consumption than their matched controls in the three months following the outbreak and 8.6% over the two years through 2021, roughly 2,900 yuan per year of platform-recorded spending for the average consumer (Table 4). The response operates almost entirely through higher spending conditional on transacting rather than through more frequent platform use, and it is largest where a liquidity backstop should have the highest marginal value: among low-spending consumers, in less developed regions, among consumers without other active credit instruments, and among recent entrants to digital payments. Yet the gain involves almost no debt. Over the two post-shock years, 86% of treated consumers never carry an interest-bearing balance; in an average month, 0.29% of them are overdue, peaking at 0.52% at the height of the shock. These are the paper’s headline answers: the credit line functioned as a consumption buffer for the consumers who needed it most, and it did so without the debt accumulation that motivates the regulatory concern.

These comparisons identify the effect of access under a conditional parallel-trends assumption: absent the credit line, matched users and non-users would have responded to the shock in parallel. Three features of the design support this assumption. First, the matched groups are balanced not only on the matching variables but on held-out covariates, including spending across 42 granular merchant categories (Appendix B). Second, pre-shock consumption paths are statistically indistinguishable within the matching window: quarterly event-study coefficients are individually and jointly insignificant (Wald $p = 0.42$), and sensitivity bounds following Rambachan and Roth (2023) show that the acute-period effect survives violations of parallel trends larger than any observed pre-shock deviation (Appendix D.2). Third, the baseline is, if anything, conservative: adjusting for control consumers who adopt the line after the shock and treated consumers who leave the platform raises the estimates, estimators that relax the matching (inverse-probability weighting, wider calipers, reweighting) all produce larger effects, and the matched sample’s per-capita consumption tracks the working-age-adjusted national benchmark (Section 3). The residual concern, selection on unobservables that begins to matter exactly at the shock, cannot be excluded by design in any study of this kind; Section 1.5 discusses it, and the mechanism evidence described next bears on it directly, because the alternative interpretations it rules out are the ones selection would generate.

To interpret the combination in the headline results, a large and persistent consumption response with almost no borrowing, we turn to the theory of household liquidity. In the buffer-stock model of Carroll (1997), consumers facing uninsurable risk hold precautionary balances and forgo consumption to maintain them. A committed credit line substitutes for those balances: it truncates the worst outcomes whether or not it is ever used, so a consumer who gains access holds a smaller buffer and consumes more in every period, while the line itself is drawn only in the states where bad shocks realize. Availability, not borrowing, does the work. This channel has

been hard to isolate, because in most settings a credit expansion is observed through the borrowing it generates, and its effect is measured the same way (Gross and Souleles, 2002; Aydin, 2022). Our data record access and use separately for nearly half a million consumers, so the borrowing and availability channels have different observable implications.

Three sets of facts, developed in Section 4, line up with the availability channel and against its alternatives. First, the response rises monotonically in every proxy for the tightness of the liquidity constraint, with no reversal on any margin: over two years, the effect rises from 5.0% in the top decile of pre-shock spending to 16.2% in the bottom decile; from 5.9% for consumers who used all four payment instruments to 27.3% for those who used only one; from 6.2% for active credit-card users to 19.2% for consumers without active card use; and from 3.1% for the oldest digital-payment cohort to 17.2% for the newest (Figures 4 and 5). Second, the persistent part of the response sits exactly where a relaxed precautionary motive predicts. The effect on daily necessities grows from 7.7% in the lockdown quarter to 15.1% in 2021, by then the largest of any category, while big-ticket and travel effects fade to zero and committed payments, a placebo category that a credit line cannot plausibly affect, do not respond in the acute period. The response is concentrated in in-person rather than online spending, and it survives on the intensive margin of categories the consumer already transacted before the shock, where payment-channel migration is impossible by construction (Table 8). Third, and decisive for separating availability from borrowing, credit-line usage is ordered inversely to the consumption response: single-instrument consumers, whose response is 27%, use 2.8% of their line on average, and 93% of them never carry an interest-bearing balance, while four-instrument consumers, whose response is 6%, use 7.0% of theirs (Figure 7). Consumption rises most where the line is touched least. Promotional targeting, wealth effects, platform engagement, and borrow-and-spend accounts each fail at least one of these facts; a latent liquidity buffer fits all three.

The two regulatory concerns are not a trade-off in this setting: the consumption support is real and concentrated among consumers with the fewest alternatives, and the debt accumulation that would normally be its price does not materialize, because the channel does not require borrowing. Three boundary conditions qualify the answer. The effect is attenuated in cities under the strictest containment measures; a credit line relaxes the household's liquidity constraint but cannot offset closed spending opportunities. A minority of users, 14%, do carry interest-bearing balances at some point, and delinquency, while low, is not zero. And repayment discipline is supported by the platform's integrated design, from automatic deduction to the value consumers attach to continued access, so the financial-health results need not extend to unbundled buy-now-pay-later products in markets with weaker repayment infrastructure. Within these bounds, the evidence characterizes payment-embedded credit as liquidity insurance for thin-file consumers rather than as a debt product.

Our paper contributes to three strands of literature. The first is the literature on fintech and *financial inclusion*. New lending models built on alternative data and automated screening (Berg, Burg, Gombović, et al., 2020; Vissing-Jorgensen, 2011; Lee, Yang, and Anderson, 2025; Ghosh, Vallée, and Zeng, 2026) have extended credit to consumers and firms historically outside the formal financial system, especially in developing economies (Frost et al., 2019; Agarwal, Alok, Ghosh, and Gupta, 2019; Di Maggio and Ratnadiwakara, 2026; Liu, Lu, and Xiong, 2022), with documented effects on firm growth and volatility (Hau et al., 2024; Eça et al., 2022; Chen, Huang, et al., 2022).³ The payment technology in which such credit is embedded is itself a force for inclusion: debit cards, mobile money, and cashless payments raise saving, risk sharing, and spending in developing economies (Jack and Suri, 2014; Suri and Jack, 2016; Bachas et al., 2021; Crouzet, Gupta, and Mezzanotti, 2023; Agarwal, Ghosh, et al., 2024; Higgins, 2024), and within the Alipay ecosystem payment adoption itself expands access to platform credit (Ouyang, 2026). Evidence on the household side of big tech credit, by contrast, remains limited to early experiments, short post-adoption windows, or outcomes other than consumption (Bharadwaj, Jack, and Suri, 2021; Ji et al., 2023; Bian, Cong, and Ji, 2024; Dong et al., 2025). We quantify the consumption response to a credit line used by half a billion consumers, most of whom hold no traditional credit card, tracking matched users for two full years around a common shock and observing credit access and credit use separately.

The second is the household-finance literature on *consumption responses to shocks* and the role of credit (see Jappelli and Pistaferri (2010) for the framework and early evidence; institutions studied include asset buffers, informal risk sharing, and unemployment insurance, Rosenzweig and Wolpin (1993), Townsend (1994), Gruber (1997)). Theory attributes the transmission of shocks to consumption to liquidity: consumers who face borrowing limits accumulate precautionary buffers (Zeldes, 1989; Deaton, 1991; Carroll, 1997), and consumption responses track liquid assets rather than net worth (Kaplan and Violante, 2014). Whether credit access shapes the consumption response to shocks is unsettled: Baker (2018), Agarwal and Qian (2014), Baker and Yannelis (2017), and Gelman, Kariv, et al. (2020) find that credit and liquidity constraints drive spending responses, while Ganong and Noel (2019), Baugh et al. (2021), and Hundtofte, Olafsson, and Pagel (2019) find little use of credit to smooth income shocks.

Our contribution to this debate is a distinction the existing evidence could not draw. Studies of credit expansions measure responses that operate through borrowing (Gross and Souleles, 2002; Agarwal, Chomsisengphet, et al., 2018; Aydin, 2022); buffer-stock theory (Carroll, 1997) predicts that a committed credit line should also raise consumption through its mere availabil-

³An adjacent literature studies the supply side of fintech lending in developed markets, including marketplace lending (Tang, 2019; Vallée and Zeng, 2019) and technology-based mortgage origination (Buchak et al., 2018; Fuster et al., 2019); see Berg, Fuster, and Puri (2022) for a survey.

ity, by substituting for precautionary balances, and credit lines can insure households against shocks precisely because their terms do not reprice when the shock hits (Braxton, Herkenhoff, and Phillips, 2024).⁴ Because we observe credit access and credit use separately for half a million consumers, we provide direct evidence on this availability channel at scale: the consumption response is largest precisely among consumers who never draw the line. Recent studies document spending responses to new cards or higher limits that likewise do not operate through borrowing, attributing them to mental accounting (Gelman and Roussanov, 2024), consumption cues (Agarwal, Alok, Ghosh, and Zhang, 2024), or income inference (Yin, 2026); in our setting, the constraint gradients, the necessity tilt of the response, and the unresponsive committed payments point instead to a precautionary buffer, with the response largest exactly where the shadow value of liquidity is highest.

The third is the emerging evidence on *buy-now-pay-later* credit and on household spending through the pandemic. BNPL access raises consumer spending and merchant sales (Di Maggio, Williams, and Katz, 2022; Bian, Cong, and Ji, 2024; Berg, Burg, Keil, et al., 2025) but can lead to over-borrowing (deHaan et al., 2024); in China, an early experiment with the product we study finds that access raises consumption even among users who do not borrow (Ji et al., 2023). Pandemic-period studies document the collapse and recovery of consumer spending in the United States and China (Baker, Farrokhnia, et al., 2020; Kim, Parker, and Schoar, 2025; Chen, Qian, and Wen, 2021; Parker et al., 2022; Chetty et al., 2024), the effects of China’s digital stimulus coupons (Liu, Shen, et al., 2021; Ding et al., 2025), and shifts in credit use on the same platform through the pandemic (Hu, Huang, and Liu, 2024).⁵ Relative to this work, we track matched consumers for two full years after a common liquidity shock and show that a BNPL-type credit line supported consumption persistently, with low delinquency and without widespread revolving debt, and with the response concentrated among consumers who never borrow.

The remainder of the paper proceeds as follows. Section 1 describes the institutional setting, the data, the matching design, and identification. Section 2 reports the baseline consumption effects and their composition, and Section 3 presents robustness and external validity. Section 4 develops the buffer framework and tests its predictions. Section 5 examines delinquency, revolving behavior, and the distribution of the benefits across demographic groups; Appendix F reports complementary city-level evidence.

⁴Related evidence on the insurance value of credit access includes Herkenhoff (2019) on consumer credit and unemployment spells and Jaffe, Malani, and Reif (2026), who show that access to credit lowers households’ valuation of formal insurance.

⁵Cox et al. (2020) document how pandemic spending declines varied with household liquidity in the United States.

1 Data and Methodology

Given the critical dependence of our methodology on the unique aspects of our data source from Alipay, we will start this section by detailing the users and services of Alipay and Huabei. Then, we will explain how we design our empirical strategy based on these data characteristics, followed by the details of the sampling procedure, the variables, and the specifics of the matching method.

1.1 Alipay and Huabei

Alipay is a payment platform that provides customers and merchants with online and mobile payment services. As of June 2020, Alipay is used by 1.3 billion users and 80 million merchants. It was created in 2004 by Alibaba to provide a trustworthy third-party payment method for e-commerce buyers and sellers. With the adoption of QR code payment in 2014, it has quickly expanded from online e-commerce to offline face-to-face use cases,⁶ becoming a universally accepted payment tool in China. Alipay is the largest payment platform in the mobile payment market in China, with about 45% market share, followed by WeChat Pay and UnionPay.⁷

Huabei, meaning *just spend* in Chinese, is a consumption credit service that Alipay has offered to its Chinese customers since December 2014. Although without a physical card, consumers can use Huabei as a payment method in their digital wallet. The usage is similar to a typical credit card. Within a monthly billing cycle, consumers can spend as they wish within their credit limit. After the closing date, they can choose to pay everything back before the due date to avoid interest or make install payments over 3 to 12 months with an annual interest rate (APR) of approximately 15%.⁸ Furthermore, it was also one of the earliest to offer the “buy now, pay later” option, typically used for large purchases, allowing consumers to directly select the payment period at the time of payment. By June 2020, Huabei was used by about 500 million users, with an average outstanding balance of 2,000 *yuan*, and the credit line was as low as 20 *yuan*.⁹

Our analysis makes use of the de-identified payment data provided by Luohan Academy and Ant Group, the parent company controlling both Alipay and the credit service Huabei. Generated

⁶This expansion particularly accelerated after 2016, when the People’s Bank of China formally introduced industry standards for QR code payments. During the period of our study, from 2017/01 to 2021/12, we observed a steady rise in the average monthly expenditure for face-to-face payments before 2019. However, the growth in online payments remained relatively stable throughout this period. The proportion of expenditure not related to e-commerce steadied at around 75% after 2019.

⁷According to multiple reports of iResearch and Statista, Alipay’s market share in China’s mobile payment market decreased from approximately 80% in 2014 to about 50% in 2022, but it has remained the largest payment platform.

⁸Typically, the bill closing date is the 1st of the following month, and the payment is due by the 9th of that month, resulting in a maximum interest-free period of 40 days. Choices for instalment payments include durations of 3, 6, 9, and 12 months, with Annual Percentage Rates (APR) of 14.94%, 15.27%, 15.34%, and 15.96%, respectively. It was the largest consumer credit service in China both in credit balance and number of consumers.

⁹See Ant Group’s prospectus of its proposed IPO at the Hong Kong Stock Exchange on October 26, 2020: <https://www1.hkexnews.hk/listedco/listconews/sehk/2020/1026/2020102600165.pdf>

from transaction-level data, our final dataset provides monthly information on consumers' payments between January 2017 and December 2021. We also use individual information, including age, gender, education, residential cities, etc. We will first explain the empirical strategy and then get into the details of the data and our sample.

1.2 Empirical strategy

We use the lockdown caused by the COVID-19 pandemic in 2020H1 as a negative impact on consumer income to study how access to fintech credit can help alleviate the income shock. In the several months following the lockdown of Wuhan on January 23, 2020, almost all the cities in China experienced different levels of lockdowns. The lockdown and mobility restriction measures were gradually relaxed from late February to June. As a result, the total consumption decreased by 19% and 3.9% in the first and second quarters of 2020 on a year-over-year basis.¹⁰

The main challenge for our identification is that having access to Huabei and the consumption after shock may be affected by unobserved confounding factors. We fully exploited the two advantages of our data source to resolve this problem.

The first feature of the data source is its large user base, which allows us to select a considerably large sample even after applying restrictive conditions. One hypothesis is that consumers who started using Alipay earlier or use it more frequently are more likely to utilize the credit service. Those who have not used the credit service tend to be inactive on the digital payment platform. The large user base allows us to impose demanding activeness criteria, six consecutive months of digital-payment usage without any use of the credit service, and still retain millions of candidate control consumers (Section 1.3).

The second is that our data includes rich individual characteristics that are of economic significance. We have accurate information on gender, age, education, birth city, residential city, and the time to start using digital payment, and a rich set of information predicted by algorithms such as occupation, car ownership, and real property ownership. Therefore, we can construct a control group that is almost identical to the treatment group *ex-ante*, and use additional variables to verify the quality of the match.

These two features allow us to construct the control group with the matching method. Having constructed the control and treatment groups, we estimate the treatment effect using Poisson Pseudo-Maximum Likelihood (PPML) as our primary specification:

¹⁰Using high-frequency daily transaction data, Chen, Qian, and Wen (2021) find that the consumption decreased by 32% in the 12 weeks following the lockdown. According to China's National Bureau of Statistics, the total retail sales of consumer goods decreased by 19% and 3.9% in the first and second quarters of 2020 on a year-over-year basis (see https://www.gov.cn/xinwen/2020-07/16/content_5527279.htm).

$$Y_{it} = \exp(\beta(\text{Treatment}_i \times \text{Post}_t) + \lambda_i + \delta_{ct}) \epsilon_{it} \quad (1)$$

where Y_{it} is the consumption amount for consumer i in month t , λ_i are individual fixed effects, and δ_{ct} are city-by-month fixed effects. PPML naturally handles zero consumption observations and provides consistent estimates of semi-elasticities without requiring log-like transformations (Silva and Tenreyro, 2006; Chen and Roth, 2024). The coefficient β represents the treatment effect as a semi-elasticity: the percentage change in consumption attributable to fintech credit access.

Throughout the paper, the pre-shock estimation window is December 2018 through December 2019, the 13 months preceding the January 2020 outbreak; this is also the window over which the matching variables are defined. Post-shock horizons are the acute three-month window (February–April 2020, aligned to the February trough of the shock), the first half of 2020, calendar year 2020, and the full two years 2020–2021. Earlier months (2017–2018) enter only descriptive and placebo exhibits and are not part of the estimation window.

1.3 Data

We construct our sample from the Alipay consumers who are active in each of the six months in the second half of 2019 (henceforth 2019H2). We first randomly selected 1,000,000 consumers who had used Huabei at least once in 2019H2 (treatment) and 5,000,000 consumers who had not used Huabei before the end of 2019 (control). Then, we exclude those individuals with excessively high or low consumption, lack individual information, and individuals who are not within the age range of 18 to 60 years, as Huabei is unable to offer services to them¹¹. Among these consumers, we then randomly draw 500,000 and 3,000,000, respectively, as the treatment and control samples.

Variables. The data we use is from January 2017, shortly after Alipay widely promoted its QR code payment and went from online payment into a widely accepted payment in both online and face-to-face transactions, and ended in December 2021, two years after the COVID-19 shock. The variables include (1) consumers’ payment information with Alipay, including payment methods, merchant types, and the respective amount and number of payments at the monthly level; (2) individual characteristics, including age, gender, education, occupation, residential cities, the date starting using digital payment; and (3) balance and gains of investment in mutual funds and money market funds on Alipay’s wealth management platform.

Table 1 shows the summary statistics of the original sample of 3,500,000 consumers, including age, gender, the quarter in which consumers started using Alipay, and the consumption of

¹¹In the initial random sample, 94% of consumers are matched with complete individual information. We drop consumers (1) whose age was below 18 or above 60 in 2019; (2) whose total consumption in 2019H2 is at the bottom 5% (with a threshold of *yuan* 1,005); (3) whose monthly consumption is over *yuan* 100,000. This leaves 882,618 treated and 3,063,787 control consumers.

each year. In 2019, individuals in the treatment and control groups spent an average of 31,888 *yuan* and 22,222 *yuan*, respectively, magnitudes comparable to official per-capita consumption expenditure; a formal benchmarking exercise follows below. The statistics also reveal clear differences between the two groups in both composition and consumption levels, which motivates the matching procedure in Section 1.4.

Table 3 benchmarks the matched sample’s annual consumption against external sources: the National Bureau of Statistics per-capita consumption expenditure for 2017–2021 and, for 2018, the China Family Panel Studies survey. Adjusting the benchmarks by the population share aged 18–59 to match our sample frame, sample consumption ranges from 90% to 126% of the national average and 72% to 99% of the urban average over 2018–2021; 2017 is the exception, when digital payments were still spreading. The 2018–2019 median-to-mean ratio of 0.76 is likewise in line with the national income distribution. The spending recorded on the platform therefore captures a representative share of total consumption for these consumers.

Supplemental data. In addition to our primary data set, we incorporated several public datasets for further analysis. (1) We obtained the asset balance and gains on two types of assets, mutual fund and money market fund, of consumers in the matched sample who invest with Alipay’s wealth management platform. (2) For a sub-sample, we obtained the total credit limit from financial institutions recorded in the central bank’s credit reporting system through the payment platform. (3) We gathered information on the 2020 lockdown measures implemented in response to the pandemic from the official announcements of different provinces, with pandemic data sourced from the Chinese Center for Disease Control and Prevention public releases. (4) We use city-level information, such as population, GDP, and retail sales, from the China Statistical Yearbook and CEIC.

1.4 Matching

We use one-to-one matching to construct the control and treatment groups and then estimate the treatment effect (Abadie and Imbens, 2006; Imbens, 2015). Given the large size, long period, and rich set of characteristics of the data, we can select a number of matching variables to obtain a balanced match. Meanwhile, we deliberately hold out a set of variables to validate the matching quality. Specifically, we use the three aspects of variables to construct comparable treatment and control groups that are balanced with respect to observable characteristics.

(a) **consumption:** the logarithm of consumption in the first, second, third, and fourth quarters of 2019, as well as in December 2018. We use the quarterly consumption of 2019 to control the consumption level, hence the income level of consumers. Consumers with the same income level may consume differently between the usual time and festivals; we add the consumption in

December 2018 to control the spending pattern in the festival.

(b) **demographic information**, including *gender*, *age*, and *region*. We use gender and age to account for fundamental consumer demographics. In our dataset, the geographical information for consumers is accurate up to the city level. However, aligning solely on the city level would drastically diminish the number of successful matches. On the other hand, matching at the broader provincial level is overly generalized. To find a middle ground, we first identify major cities, including super cities such as Beijing and Shanghai, provincial capitals, and the leading economic cities within a province that aren't the capital, adding up to 40 cities.¹² Then, we group the areas excluding these major cities into a single region for each province. This approach leads to the definition of 70 distinct regions.

(c) **digital age**, when individuals started using digital payment, measured in quarters, *time enter*. A potential source of endogeneity in the decision to use fintech credit is consumers' familiarity with digital payments. Age, consumption level, and place of residence alone are insufficient to control consumers' payment habits in the digital economy. Furthermore, the penetration of digital payments might differ across regions, and all the transactions we observed in our study were made through digital payments. To address this issue, we control for the time when a consumer first adopted digital payments to mitigate this endogeneity. Since the widespread adoption of digital payments in China started around 2016, we do not differentiate among those who began using it before 2017.

To summarize, we have used the following nine variables in the matching process:

- consumption amount in the first, second, third, and fourth quarters of 2019, and in December 2018;
- demographic information, including *gender*, *age*, and *region*;
- *enter time*, the quarter when individuals started using digital payment.

We apply exact match on *gender*, *region*, and *enter time* and nearest match with Mahalanobis' distance for other variables with a caliper equal to 0.2.¹³ Before matching, the treatment group consists of 500,000 consumers, and the control group has 3,000,000 consumers. The pool thus offers six candidate controls per treated consumer. In total, 217,199 consumers in the treatment

¹²Specifically, the 40 big cities include 4 provincial-level municipalities: Beijing, Shanghai, Tianjin, Chongqing; the capitals of 27 provinces; 5 sub-provincial cities: Shenzhen, Ningbo, Qingdao, Dalian, Xiamen; and four non-provincial capital cities that rank first in GDP within their respective provinces: Suzhou, Quanzhou, Tangshan, and Ordos.

¹³The standard caliper used in the literature is 0.25 (Lunt, 2014). However, given our large sample size, we chose a lower caliper to further reduce the bias of the two groups. The drawback of a smaller caliper is that more subjects would not be matched. Our analysis of different calipers shows the decrease is acceptable.

group are matched with one pair in the control group, with a success rate of 43.4%. This yields a matched sample of 434,398 consumers.

We use the full matched sample of 434,398 consumers (217,199 pairs) as our baseline sample, without further filtering for platform activeness. This choice preserves the integrity of the matched sample and avoids introducing selection bias from post-matching sample restrictions. The matched sample exhibits excellent balance: the average annual consumption in 2019 is 24,888 *yuan* for the treatment group and 25,047 *yuan* for the control group, with mean ages of 30.4 and 30.5, respectively (Table 2).

Balance on held-out characteristics. Figure 2 reports the absolute standardized mean difference $|SMD|$ between the treatment and matched control groups for twelve individual characteristics that we deliberately held out of the matching, grouped into four semantic dimensions: *personal characteristics* (native-resident status, real-estate and car ownership, zodiac sign), *pre-match consumption* (total log consumption in 2017 and 2018), *digital age* (calendar year of first Alipay and first Taobao registration), and *education* (high school or below, post-secondary, college, graduate). Following Imbens and Rubin (2015), we benchmark against the conventional $|SMD| < 0.10$ threshold. Nine of the twelve covariates lie below 0.10; the remaining three (Alipay registration year, Taobao registration year, and college indicator) sit just at or slightly above 0.10, well within the looser 0.25 threshold. Two pieces of evidence are particularly informative. First, total log consumption in 2017 (a year entirely excluded from the matching procedure) has $|SMD| = 0.001$, providing a clean placebo confirming that the matching does not generate spurious balance only on the matched periods. Second, the zodiac-sign indicator, mechanically determined by birth date and orthogonal to all matching variables, has a maximum $|SMD|$ across twelve signs of 0.013, suggesting that the matching captures all dimensions of consumer heterogeneity that systematically correlate with Huabei adoption, including the latent dimensions that birth date proxies for.

Balance across granular spending categories. To assess whether the match quality extends beyond aggregate covariates, we compute normalized differences (Imbens and Wooldridge, 2009) between treatment and control groups across 42 merchant category codes (MCCs) spanning six economically meaningful dimensions: risk and behavioral preferences, income-sensitive spending, inelastic demand, digital engagement, non-durable consumption, and durable goods. Following Imbens and Wooldridge (2009), we use $\Delta_{ND} = (\bar{X}_T - \bar{X}_C) / \sqrt{S_T^2 + S_C^2}$, which is invariant to sample size. On the unconditional mean, the dimension that directly corresponds to our DiD outcome, *all 42 MCCs have* $|\Delta_{ND}| < 0.05$, with a maximum of 0.043 (department stores) and a median of 0.009, far below the stringent threshold of 0.10 used in experimental studies. Appendix B

reports the full results by category along with intensive-margin breakdowns. This confirms that the matched pairs are balanced not only in aggregate covariates but also in the composition of spending across granular merchant categories.

1.5 Identification: Strategy and Threats

Our identifying assumption is *conditional parallel trends*: absent access to fintech credit, the post-shock consumption of treated and matched control consumers would have evolved in parallel, conditional on the pre-shock covariates used in matching. The COVID-19 outbreak in January 2020 serves as a common, quasi-exogenous shock to household liquidity; it is not a shock to Huabei access itself. The interaction of this shock with treatment status identifies the average treatment effect on the treated (ATT). We assess this assumption against five concerns.

Selection on unobservables. Matching balances observable pre-shock characteristics but cannot balance unobservables such as latent digital engagement, marginal propensity to consume, or optimism, which may correlate with both Huabei adoption and the post-shock consumption response. We mitigate this in three ways. First, pre-trends are flat within the identification window (December 2018 through December 2019): quarterly event-study coefficients are individually and jointly indistinguishable from zero (Wald $p = 0.42$), and the estimated linear differential trend is 0.10% per month ($p = 0.35$). Quarters in 2018, before the matched window, show small deviations of 2 to 5%, reflecting pre-match level differences; sensitivity bounds that treat these deviations as the benchmark for possible post-period violations are reported in Appendix D.2 (Rambachan and Roth, 2023), and the acute-quarter effect survives violations up to 1.5 times the largest pre-period deviation. Second, the matched groups are balanced not only on aggregate covariates but across 42 granular merchant categories (Appendix B). Third, the IPW estimator on the full sample (Section 3.3) does not rely on exact matching. We nonetheless regard a shock-triggered divergence in unobserved consumer types as the most material residual concern, and discuss it further in the Conclusion.

Anticipation and reverse causality. If consumers adopted Huabei in anticipation of the pandemic, treatment status would be endogenous to the shock. We restrict the treatment group to users who activated Huabei before January 2020, prior to any public information about COVID-19, and the flat pre-shock event-study coefficients provide no evidence of anticipatory adjustment.

Differential shock exposure. Chinese cities experienced lockdowns of differing intensity and timing, and Huabei penetration may correlate with local exposure. City $\times$ month fixed effects absorb all city-level differential shock exposure.

Measurement error from single-platform data. We observe payments only through Alipay, not WeChat Pay, cash, or bank cards. Two features limit the resulting bias. First, our sample-selection criteria require high pre-shock Alipay activity (Section 1.3), so the platform captures the bulk of these consumers’ spending. Second, accounting for treated users who subsequently reduce or discontinue platform usage and control users who later adopt Huabei yields *larger* estimates (Section 3), indicating that platform attrition biases the baseline *against* finding an effect rather than spuriously generating one.

External validity of the matched sample. Matching retains 43.4% of treated users, raising the concern that estimates do not generalize. Section 3.3 addresses this directly through IPW, relaxed-caliper PSM, and reweighting, all of which reproduce the baseline result. As further evidence on representativeness, Figure 1 shows that the per-capita consumption of the matched sample is highly consistent with the working-age-adjusted national average over 2018–2021.¹⁴

Of these, selection on unobservables is the only concern we cannot fully rule out with the current design; the remaining four are addressed by the research design or by the robustness and external-validity analyses that follow.

2 Main Results

Next, we report the baseline estimation of the treatment effect in the matched sample. We then examine the margins and channels of the response and decompose the effect into consumption categories.

2.1 Baseline estimation

Our main findings are conveyed by Table 4, which presents the PPML estimates of the treatment effect of fintech credit access on total consumption. The treatment group exhibits 9.0% higher consumption in the three months following the shock (February–April 2020), 7.8% in the first half of 2020, 8.5% in the full year of 2020, and 8.6% over the two-year period ending in 2021. All estimates are highly significant ($t > 8$) with standard errors clustered at the region level (66 regions).

Event-Study (Dynamic DID) Estimates

Figure 3 plots the event-study estimates using monthly data from 2019–2021. Pre-shock coefficients within the identification window are small and show no trend: quarterly pre-period coeffi-

¹⁴The two series diverge in 2017, when digital payment was still undergoing rapid adoption in China.

cients are jointly indistinguishable from zero (Wald $p = 0.42$), and the estimated linear differential trend is 0.10% per month ($p = 0.35$). Monthly coefficients fluctuate about 1.7% below the December 2019 reference month, a treated-specific December seasonal rather than a differential trend. Appendix D.2 reports formal sensitivity bounds. This supports the identifying assumption that matching eliminated pre-existing differences in consumption trends between the treatment and control groups.

Immediately following the shock (2020/02), we observe a sharp spike in the treatment effect. This captures the essence of consumption smoothing: as the control group was forced to cut spending due to liquidity constraints, the treatment group utilized fintech credit to buffer the shock. The effect does not revert to zero after the initial lockdown phase. It stabilizes at a positive level throughout 2020 and 2021.

2.2 Margins and Channels of the Response

We decompose the total PPML effect into an extensive and an intensive component. The extensive margin is estimated using a linear probability model (LPM) with a binary indicator for active platform usage as the dependent variable: access to fintech credit increases the probability of any positive consumption by 1.4 percentage points over two years ($t = 8.2$). The intensive margin is estimated by a log model restricted to observations with positive consumption ($Y_{it} > 0$): conditional on transacting, users with fintech credit increase their total consumption amount by 12.7% ($t = 13.5$) over the two-year period. The intensive margin accounts for the bulk of the total effect.¹⁵

Panels C and D of Table 4 locate the response. The effect is concentrated in in-person spending, 12.9% over two years, while online spending is essentially unchanged at every horizon beyond the acute window (Panel C). Within in-person spending, both the number of payments (5.3% over two years) and the amount per payment (7.2%) contribute to the response (Panel D). The concentration in in-person consumption is informative about the mechanism, and Section 4.3 returns to it.

Appendix E reports the full decomposition of the consumption response across online and in-person channels, transaction frequency, and the extensive and intensive margins, for both the 434K matched sample and the externally reweighted sample.

¹⁵Because the conditional-on-positive sample $\{Y_{it} > 0\}$ is itself affected by treatment through the extensive margin, the intensive-margin coefficient confounds a behavioral effect with a composition effect and is not a causally identified parameter (Chen and Roth, 2024). We therefore report the extensive/intensive split only as a descriptive decomposition; all causal inference relies on the PPML estimate in levels, which requires no conditioning on positive consumption. For the same reason we do not report a dynamic (event-study) version of the intensive margin.

2.3 Decomposition of Consumption Categories

For the consumption payment not on the e-commerce platform, we have the information of the Merchant Category Code (MCC) of the payment. Therefore, we can further decompose the treatment effect in different types of consumption. Following the method in [Ganong and Noel \(2019\)](#) and [Lusardi \(1996\)](#), we use the MCC data to categorize spending into 19 mutually exclusive and comprehensively exhaustive categories, and then summarize as Durable, Non-Durable, and Strictly Non-Durable expenditures. Appendix Table [A10](#) presents the treatment effect in different time horizons for each category with its respective share in the total expenditure, ranked from highest to lowest based on the short-term treatment effect, and Appendix Figure [A10](#) plots the corresponding monthly dynamic effects by durability and necessity groups. The effect is broad-based across category groups: the restaurant and discretionary responses are front-loaded at the shock, while the necessity response builds gradually over the two-year horizon. This durability taxonomy describes the shape of the response; Section [4.3](#) regroups the same merchant categories by the adjustability of the expenditure to test the mechanism predictions.

3 Robustness Analysis

Next, we discuss the robustness of the baseline estimation. We account for the post-shock adoption and exit of fintech credit, re-estimate the effect in the subsample with observable external bank credit, and assess the external validity of the matched sample.

3.1 Adoption and Exit

One concern with our baseline estimate is that we only controlled for whether the two groups of consumers had access to fintech credit at the start of the shock. However, after the shock, it's possible that consumers in the control group could start using fintech credit, while those in the treatment group might stop using it. Table [A14](#) shows that accounting for adoption and exit yields larger estimates, confirming that our baseline results are conservative.

Some control-group consumers began using fintech credit after the shock. By the end of 2021, almost one-third of control-group consumers had adopted fintech credit, and the activeness of these new users in using fintech credit is steady.

Given that fintech credit usage might increase consumer spending, considering control group users who adopt fintech credit as treated would likely increase our estimate of the treatment effect. To assess this impact, we consider control group consumers who adopt fintech credit as treated, starting from the month they began using fintech credit. Column (2) in Table [A14](#) shows a long-term treatment effect of 13.6% under such specification, higher than the baseline

effect. Appendix Figure A9 reports the corresponding monthly dynamic effects: panel (a) plots the baseline and adoption-adjusted series for total consumption together with the monthly count of new control-group adopters, while panel (b) decomposes the effect into in-person and online (e-commerce) channels. The adoption-adjusted estimates sit above the baseline throughout the post-shock window, and both series highlight that the effect is concentrated in in-person consumption, a pattern Section 4.3 returns to.

Meanwhile, some consumers in the treatment group may stop using fintech after the shock. This could lead to an underestimation of the actual treatment effect in our baseline estimate. It should be noted, however, that a consumer’s non-use of credit does not necessarily mean that the credit limit has no value to them. To examine the impact of exits on our estimates, we defined an exit indicator: if a consumer stopped using fintech credit from a given month through to the end of our data in 2021/12, the consumer is considered as having exited fintech credit for that month. Column (3) in Table A14 shows the long-term treatment effect is 17.8% after considering the exit, higher than the baseline.

In addition, in column (4), we consider both the adoption and exit, providing a more precise control over the effect of using fintech credit. The treatment effect is 20.0%. This further supports our claim that the baseline estimate does not overstate the effect.

3.2 External-Credit-Observed Subsample

Another concern is that the matched-sample result might be driven by consumers without observable formal-credit relationships. We therefore re-estimate the baseline effect in the subsample of matched consumers with observable external bank credit lines. This is a more financially connected subsample and should, if anything, have lower marginal value of an additional Huabei line.

Table A15 shows that the treatment effect remains positive and statistically significant in this subsample. The implied effect is 12.0% in the three-month window and 6.1% over two years. The smaller long-run estimate is consistent with the buffer interpretation: consumers with observable formal-credit relationships are less constrained on average, so an additional fintech credit line has lower marginal value. One caveat applies: credit-bureau coverage is mechanically related to treatment status, since Huabei is itself a credit relationship (Section 4.2); we therefore read this subsample check as corroborating, rather than independent, evidence.

3.3 External validity of the matched sample

The external validity of our matched-sample estimates warrants attention: the 1:1 matching procedure retains 217,199 treated users, a 43.4% match rate, and if these matched users differ system-

atically from the unmatched majority, the estimated ATT may not generalize to the full Huabei population. We address this through three complementary estimation specifications that progressively relax the matching restriction: (i) inverse probability weighting (IPW) estimation on the full 3.5-million-consumer sample, imposing no matching restriction; (ii) many-to-one propensity-score matching with varying K and calipers, which achieves up to 100% match rates; and (iii) reweighting the 434K matched sample by inverse probability of selection to restore representativeness of the full treatment population (Stuart et al., 2011; Hartman et al., 2015). Table 5 reports the PPML treatment effect across all methods and four post-windows (3 months, 6 months, Year 2020, 2 years).

Three patterns emerge. First, every estimate is positive and highly significant, and the matched baseline is the smallest across methods at every horizon, so any external-validity bias pushes the baseline *down*, not up. Second, the 1:K PSM estimates on the broader samples are substantially larger (19–21% at two years, virtually identical across the three calipers) than the matched baseline (8.6%); the IPW estimate on the full 3.5M panel is intermediate (14.8%); and the reweighted matched sample (8.6%) is essentially identical to the unweighted matched baseline. The ordering is economically interpretable: the matched baseline excludes the $\sim 283,000$ treated consumers whose spending profiles are hardest to match, namely those with volatile or extreme consumption trajectories, and these are precisely the consumers for whom fintech credit delivers the largest consumption smoothing. Third, the within-method post-window pattern is stable: the three-month acute window is the largest for every row, and the 1-year and 2-year horizons settle close together, consistent with the dynamic event-study profile in Figure 3. Taken together, the broader-sample specifications indicate that our matched baseline is conservative, and that the true average treatment effect for the broader Huabei user population is at least as large as what we report.

4 Mechanisms: Fintech Credit as a Latent Liquidity Buffer

The results so far establish that access to Huabei raises consumption and that the effect is persistent. This section asks why. We first develop a simple framework in which a credit line affects consumption through its availability rather than through borrowing, and derive five predictions that distinguish this channel from its alternatives. We then test each prediction in turn, exploiting three features of the data: independent measures of how liquidity-constrained each consumer is, the merchant-category composition of spending, and complete records of how treated consumers actually use the credit line.

4.1 Conceptual Framework and Testable Predictions

Huabei extends a revolving credit line, not a loan. At adoption the consumer receives a credit limit; whether and how much to draw is decided transaction by transaction at the point of sale, repayment within the monthly billing cycle is interest-free, and only balances rolled over into installments accrue interest. Two features of this contract matter for interpreting our results. First, access and use are distinct: a consumer can hold the line for years without ever carrying a balance. Second, because the line is embedded in the payment instrument itself, drawing on it involves no application, delay, or transaction cost. Huabei therefore provides its holder with a guarantee: whatever happens, up to the credit limit, next month’s consumption need not fall below the limit.

The natural framework for evaluating such a guarantee is the buffer-stock model (Deaton, 1991; Carroll, 1997). Consumers who face uninsurable income risk and cannot borrow hold a stock of liquid assets as self-insurance, and they forgo current consumption to build and maintain that stock. The target buffer, and hence the consumption sacrificed to maintain it, is largest for consumers with volatile income and no access to formal credit. A committed credit line is a close substitute for buffer assets: it truncates the worst outcomes whether or not it is ever used. A consumer who gains access re-optimizes by holding a smaller precautionary buffer and consuming more in every period. Crucially, the line is actually drawn only in the states of the world in which bad shocks realize; in all other states consumption rises while the line sits idle. Availability, not borrowing, does the work.

This availability channel is distinct from the mechanism documented in the existing consumer-credit literature. Gross and Souleles (2002), Agarwal, Chomsisengphet, et al. (2018), and Aydin (2022) show that consumers respond to exogenous credit expansions and that constrained borrowers respond most, but in these settings the response is measured as, and operates through, an increase in borrowing. Under a pure borrowing channel, consumption rises because balances rise. Under the availability channel, consumption rises even where balances do not. Credit limits function as part of the household’s buffer stock (Fulford and Schuh, 2024), and the corporate-finance literature has long treated committed credit lines as liquidity insurance that is valuable even when undrawn (Holmström and Tirole, 1998; Sufi, 2009; Acharya et al., 2014). Our data record access and use separately, so the two channels have different observable implications, and the alternatives to a liquidity mechanism (a wealth effect of perceived new resources, promotional subsidies tied to adoption, temporary pandemic relief, and mechanical growth of platform-recorded spending) have different implications still.

We organize the analysis around five predictions.

P1. *The effect increases in the tightness of the liquidity constraint.* The shadow value of a liquid-

ity guarantee is highest for consumers with the fewest existing substitutes: low income, no credit card, few payment instruments, short credit history. Any liquidity-based mechanism implies this gradient; a promotional or platform-engagement story does not, because promotions target high-value, not low-liquidity, consumers.

- P2. *The permanent component of the effect is concentrated in routine, recurring consumption.* The precautionary motive binds on the consumption flow that must be financed every month: groceries, daily services. Once the buffer motive is relaxed, this flow rises and stays higher. Postponable expenditures respond differently: a durable purchase brought forward produces a temporary spike, not a permanent shift. A wealth effect, by contrast, would raise spending roughly proportionally across categories.
- P3. *The response does not require borrowing.* This is the prediction that separates availability from the borrowing channel. If consumption is financed by drawing the line, usage must rise with the response; under the availability channel the ordering can be inverted, because the most constrained consumers hold the line as insurance precisely in order not to need it.
- P4. *Committed payments do not respond.* Expenditures that the consumer cannot adjust (property-management fees, transit fares, taxes) are unaffected by a relaxed liquidity constraint. This placebo also bounds a mechanical alternative: if measured effects reflected growing use of the platform rather than consumption, these categories would grow like any other.
- P5. *The effect persists after the pandemic.* The line remains available once the lockdown is lifted, whereas control consumers, having drawn down their liquid balances, have reason to hold back routine consumption to rebuild them; the gap should therefore remain positive well beyond February to April 2020. Explanations tied to the outbreak itself, such as outbreak-period promotions or repayment relief, predict instead an effect that decays toward zero by 2021.

Throughout, estimation follows the design of Section 1: PPML on the matched sample of 434,398 consumers, individual and city-by-month fixed effects, standard errors clustered by region, and the pre-shock window of December 2018 through December 2019. Reported percentage effects are $\exp(\beta) - 1$.

4.2 Heterogeneity by Liquidity Constraints

We measure the tightness of the liquidity constraint along five margins, each from a different source. Two are reported in Figure 4: the decile of pre-shock (2018–2019) total consumption, an income proxy, and the development level of the consumer’s region. Three are constructed from payment behavior and reported in Figure 5 and Table 6: the number of distinct payment instruments the consumer used in 2019 (Alipay balance, debit card, credit card, money-market wallet), a direct measure of the set of payment and credit substitutes; an indicator for active credit-card use in 2019 (at least 500 yuan of card spending), the closest analogue to the credit-access margins studied in the credit-card literature; and the platform-entry cohort, the year of first digital-payment activity, which is also an exact matching variable.

The treatment effect declines monotonically in every measure of liquidity access. Over two years, the effect falls from 16.2% in the bottom pre-shock consumption decile to 5.0% in the top decile, and from 19.3% in the least developed regions to 7.0% in the most developed (Figure 4). It falls from 27.3% for consumers who used a single payment instrument to 5.9% for those who used all four, and from 19.2% for consumers without active credit-card use to 6.2% for active card users. Across entry cohorts the effect rises from 3.1% for consumers active before 2010 to 17.2% for the 2019 cohort. The spread between the most and least constrained groups is a factor of three to five on every margin, and no measure produces a reversal.

The uniformity of these gradients is the central fact of this subsection. The five proxies are correlated, but they come from independent sources and capture different margins of access: income, local development, instruments, formal credit, and history. That the effect lines up monotonically with all five, including the measures constructed purely from payment behavior, supports P1 and is difficult to attribute to any single confound. It is also inconsistent with promotional targeting, which would concentrate effects among high-value users, the group where we estimate the smallest effects.

Bureau-recorded external credit provides a complementary but more limited view. Appendix Table A11 shows persistent effects for consumers with low recorded external credit and effects that fade by two years for consumers with high external credit. We treat this evidence as descriptive corroboration rather than as an independent test, because credit-bureau coverage is mechanically related to treatment status: Huabei itself is a credit relationship, so treated consumers are recorded at much higher rates than controls. The payment-based measures above do not suffer from this coverage problem.

4.3 Consumption Composition and Dynamics

P2 concerns what the additional consumption buys. Two facts about the location of the response set the stage. First, the effect is concentrated in in-person spending rather than online purchases (Table 4, Panel C, and Appendix Figure A9, Panel b), which is difficult to reconcile with online-promotion or checkout-salience explanations. Second, within in-person spending, the response is broad-based across merchant categories (Section 2.3). Here we sharpen the category evidence into a test of P2 and P4.

We assign in-person merchant-category codes to six groups ordered by how adjustable the underlying expenditure is: daily necessities (supermarkets, meat, fruit); health and medical; discretionary goods and services (dining, apparel, electronics, department stores); big-ticket and travel (hotels, airlines, furniture, luxury goods); committed payments, which a credit line cannot plausibly affect (property-management fees, subway fares, taxes, fines); and utility-type payments, where a shift of the payment channel onto the platform is most likely (utilities, telecom, tolls). Named categories cover 53% of in-person spending. Because the first post quarter coincides with the lockdown, pooling it with later months would conflate two different objects; we therefore estimate each group’s effect separately over four non-overlapping windows: February to April 2020, 2020 Q2, 2020 H2, and calendar 2021.

The composition of the response shifts from discretionary to necessity spending as the pandemic recedes (Figure 6, Table 7). In the lockdown quarter the largest coefficient is on discretionary categories, 26.8%. Aggregate spending in these categories collapsed for treated and control consumers alike, so this estimate measures a smaller decline among the treated, a cushioning of the crisis, not an absolute increase. From late 2020 onward the pattern is different. The effect on daily necessities rises monotonically across the four windows, from 7.7% during the lockdown to 15.1% in 2021, and by 2021 it is the largest effect of any group. The discretionary effect declines to 14.1%, and the big-ticket effect, 11.8% during the lockdown, is statistically zero in 2021. Routine consumption is the margin on which the effect is permanent; postponable consumption responds only temporarily. This is the pattern P2 predicts, and it is inconsistent with a proportional wealth effect and with a one-time stimulus, either of which would decay rather than grow in the necessity margin.

The committed-payment group supplies the placebo (P4). During the first half of 2020, when consumption categories respond most strongly, its coefficients are statistically indistinguishable from zero. A positive coefficient of 6.3% over two years emerges from late 2020, less than half the necessity effect; it is driven by property fees, which contain an income component, and by infrequent tax payments, and it is bounded by the payment-channel analysis of the next subsection. The consumption groups’ effects therefore do not reflect a mechanical expansion of platform activity.

4.4 Distinguishing Consumption from Payment-Channel Migration

The outcome variable records Alipay transactions only. Measured growth could therefore reflect a shift of existing expenditure onto the platform rather than new consumption, if treated consumers route more of their payments through Alipay as their engagement deepens. The category results already suggest this concern is limited in scope: it is the utility-type group, where re-routing a recurring bill is easiest, that shows persistent coefficients, while committed payments respond little. But the concern applies in principle to every category, so we address it directly.

The test decomposes each group’s effect into an intensive and an extensive margin. For consumers with positive spending in a category during the pre-window, the category is already transacted on the platform, and channel migration cannot generate measured growth; any treatment effect on this subsample reflects consumption. For consumers with zero pre-window spending, a treatment effect on the probability of any spending captures exactly the margin on which migration operates.

The consumption effects are intensive; the migration effects are not (Table 8). On the intensive margin the two-year estimates are essentially unchanged from the full sample: 12.4% for daily necessities, where 98% of consumers are pre-active, and 14.6% for discretionary spending. The utility-type group shows no intensive-margin effect at all (1.4%, $t < 1$): its aggregate coefficient comes entirely from consumers who begin paying bills on the platform after treatment, which is what channel migration looks like. On the extensive margin, moreover, the probability of new category spending rises more for necessities and discretionary consumption (4.8 percentage points among pre-inactive consumers) than for utility-type payments (0.95 points), the reverse of what a migration story requires. Migration exists, but it is confined to bill-payment categories and to the extensive margin, and the consumption results survive on the margin where migration is impossible by construction.

4.5 Credit-Line Usage Among Treated Consumers

Everything so far is consistent with two mechanisms: consumers borrow against the new line and spend the proceeds, or consumers spend out of resources released by a smaller precautionary buffer. P3 separates them. If borrowing finances the response, usage must be at least as large as the response, and it must be largest where the response is largest. Under the availability channel neither is true, and the ordering can invert: the most constrained consumers hold the line as insurance precisely in order not to need it. We therefore examine the credit-line records of treated consumers directly, split by the same instrument-count groups as in Section 4.2. Because comparable balance data are not observed for control consumers, this evidence is descriptive; its value is the ordering across groups.

Borrowing is rare, small, and lowest exactly where the consumption response is largest (Figure 7). Over the two post years, 86% of treated consumers never carry a positive installment balance; among single-instrument consumers, the group whose consumption response is 27%, the share is 93%. The same group uses 2.8% of its credit line on average and carries a mean balance of 403 yuan, against 7.0% and 1,604 yuan for four-instrument consumers, whose consumption response is 6%. Usage and response are inversely ordered across groups. These patterns are consistent with the full-sample revolving-balance dynamics in Figure 8, which show that persistent revolving is uncommon and short-lived. The design of the product makes this institutionally plausible: the full monthly balance is deducted automatically on the bill date unless the user actively chooses installment repayment.

This pattern excludes the borrowing channel as the driver of the average effect and matches the availability channel exactly. Consumption rises most where the line is touched least, which is possible only if the line changes behavior without being used: it substitutes for precautionary balances the consumer no longer needs to hold. Together with the necessity composition of Section 4.3, the usage evidence identifies the mechanism as a latent liquidity buffer.

4.6 Alternative Explanations

Table 9 collects the alternative explanations and the evidence on each. A borrowing channel is inconsistent with the usage patterns of Section 4.5. Payment-channel migration is bounded by the margin decomposition of Section 4.4. A wealth effect would raise spending proportionally across categories and consumers, which is inconsistent with the necessity tilt, the constraint gradients, and the unresponsive committed payments. Promotional targeting is inconsistent with the constraint gradients and with the survival of the effect when pre-period promotion exposure is controlled for. Temporary pandemic relief is inconsistent with P5: the necessity effect is largest in 2021, after the disruption ended.

Two boundary conditions on the mechanism deserve note. First, shock-intensity analyses (not reported) show that the treatment effect is, if anything, attenuated in cities subject to stricter local containment measures: credit availability relaxes the household liquidity constraint but cannot offset suppressed local spending opportunities. Second, education and the Financial Sophistication Index remain useful descriptive heterogeneity variables (Appendix Tables A12 and A13), but neither carries the mechanism: the education gradient is not monotone, and the long-run difference across FS groups is small. The evidence supports a single interpretation: access to a digital credit line raises routine consumption among consumers with few liquid substitutes by reducing their precautionary-saving motive, largely without borrowing.

5 Discussions

Two questions remain after the mechanism analysis: does the credit expansion documented above harm consumers' financial health, and who reaps the benefits? This section examines delinquency and revolving balances among fintech credit users, and then describes how the treatment effect is distributed across demographic groups. A complementary city-level analysis of fintech credit and aggregate retail sales is reported in Appendix F.

5.1 Delinquency Rate

The economic consequences of offering consumer credit to individuals are controversial, especially considering that many of the consumers using fintech credit services are relatively young, lack solid financial planning and awareness, and have lower creditworthiness. Since we only have the repayment data for the treatment group, we are unable to conduct a rigorous analysis on whether the use of fintech credit has led to excessive consumption or excessive borrowing. We will report the observed delinquency level and then compare it with similar studies.

In our sample, from 2020 to 2021, on average 0.29% of consumers in the treatment group had overdue balances each month. The monthly share ranges from 0.21% to 0.52%, with the highest value in 2020/02, the month following the shock. Among all consumers who have overdue payments, the average overdue balance is 3,458 *yuan*, while their Huabei spending bill in that month averages 8,682 *yuan*. The total overdue amount is about 0.25% of the total consumption expenditure in the two years.¹⁶

The finding is in line with a study on the consumers of the same fintech credit services by [Dong et al. \(2025\)](#), where the 30-day delinquency rate is 0.012 over a one-year period. It is also consistent with [Berg, Burg, Keil, et al. \(2025\)](#), who find that the credit losses on BNPL transactions are moderate. These contrast with [Di Maggio and Yao \(2021\)](#), who find fintech consumer credit loans are significantly more likely to be in default compared to the loans granted by traditional institutions. One non-negligible factor is that different from the fintech loan providers in many studies, the focal fintech credit service is provided by a big tech platform. Its payment system, investment platform and its own credit rating agency are all widely used by the fintech credit users. As a result, the cost of severe default might result in an exclusion from the platform, which is substantial for users, therefore leading to lower default rates, consistent with evidence and theory that platform control over valuable services can substitute for collateral in enforcing repayment ([Gertler, Green, and Wolfram, 2024](#); [Li and Pegoraro, 2026](#)).

¹⁶We should point out that, until September 2021, the regulatory framework for fintech was not clearly defined, allowing Huabei delinquencies to go unreported to the central bank's credit system. This reduced the default liability for consumers ([Dong et al., 2025](#)). Although we lack data to address the policy change on delinquency, since the policy will increase the delinquency cost, one should expect the delinquency rate to decrease after the change.

5.2 Revolving Balance and Interest Costs

A natural concern with expanding consumer credit is whether it leads to persistent indebtedness. Section 4.5 showed that borrowing is rare and lowest exactly where the consumption response is largest; here we complete the picture with the dynamics of revolving balances, defined as balances carried beyond the monthly interest-free billing cycle, among treatment-group consumers (Figure 8).

Revolving is rare and short-lived. Even at its peak, fewer than 10% of fintech credit users carry a revolving balance in any given month, and over the entire 2017–2021 sample period, 82.2% of treatment-group consumers (178,514 out of 217,141) *never* incur a positive revolving balance (panel a); the corresponding share over the two post-shock years alone is 86% (Section 4.5), the longer window mechanically offering more opportunities to revolve. Among the minority who do revolve, balances are non-trivial, reaching an average of approximately 18,000 yuan with credit utilization around 68% by late 2021 (panel b), but they are quickly repaid: the median repayment duration is 4 months, 44.6% of revolvers clear their balance within 3 months, and 81.7% within 12 months (panel c). For most consumers, revolving balances thus represent short-lived credit use rather than chronic debt accumulation, consistent with fintech credit serving as a liquidity backstop rather than a borrowing facility.

5.3 Who Benefits: Age, Gender, and Occupation

The rich set of individual characteristics allows us to describe how the treatment effect is distributed across demographic groups. Table A16 reports the estimates by gender, age, and occupation.

Panel A shows the estimation for gender. The two-year effects are very similar for males and females, 10.4% and 10.0%, respectively. This suggests that gender differences are not central to the mechanism.

Panel B examines the effect for different age groups. The estimates are broadly similar from age 25 through 45 and somewhat smaller for consumers older than 45. This pattern is consistent with the mechanism evidence above, but it is less diagnostic than the direct liquidity and substitute-credit proxies.

Panel C shows the effect for different occupations. The information on occupation is an algorithmic prediction rather than an exact value. Nonetheless, blue-collar workers show a much larger treatment effect (17.7%) than white-collar workers (7.6%) and civil servants (5.3%). These occupational patterns are consistent with the liquidity-buffer interpretation: groups with less stable or lower cash-flow positions benefit more from an immediately available credit backstop.

6 Conclusion


This paper studies the consumption response to fintech consumer credit by analyzing the behavior of 434,398 matched consumers on China’s leading digital payment platform. Using the COVID-19 outbreak as a common liquidity shock and PPML as the primary estimation method, we find that access to fintech credit enables consumers to maintain 9.0% higher consumption in the acute phase of the shock and 8.6% over the following two years.

We interpret the findings through a precautionary liquidity-buffer mechanism. The effect is largest for consumers with weaker cash-flow positions, less active use of substitute credit, narrower observed payment toolkits, and shorter digital-payment tenure. The response is concentrated in in-person consumption and broad offline categories, and it does not require persistent revolving balances. These patterns indicate that the value of fintech credit comes not only from realized borrowing, but also from the availability of a state-contingent liquidity backstop. In this sense our evidence completes an arc in the consumer-credit literature: buffer-stock theory predicts that a committed credit line substitutes for precautionary balances ([Carroll, 1997](#)), and the empirical literature has measured consumption responses that operate through borrowing ([Gross and Souleles, 2002](#); [Agarwal, Chomsisengphet, et al., 2018](#); [Aydin, 2022](#)); we provide direct evidence, at the scale of half a million consumers, that availability alone does much of the work.

Our research also sheds light on policy discussions about whether current financial-inclusion developments benefit society. Within our two-year post-shock window, delinquency among treated users is low in absolute terms, with 0.29% of users carrying an overdue balance in an average month, and persistent revolving is rare. We caution that this may understate medium-run default risk if delinquencies are reported with a lag or if regulatory forbearance during 2020–21 partially suppressed default reporting. As BNPL-style credit lines embedded in payment instruments spread across markets and draw regulatory attention, the finding that a credit line as small as 20 *yuan* can support consumption smoothing for thin-file consumers, largely without borrowing, bears directly on the design of such regulation.


Figures

Figure 1: Quarterly Consumption of the Matched Sample and the National Benchmark (2017–2021)




Note: The figure plots quarterly per-capita consumption from 2017Q1 to 2021Q4. The solid (Treatment) and dashed (Control) lines are the matched Alipay/Huabei sample of 434,398 consumers (about 217,000 per group); each point is the sum of the three within-quarter monthly group means. The shaded area is the national benchmark: the National Bureau of Statistics (NBS) per-capita resident consumption expenditure, rescaled by the 16–59 working-age population share (0.62–0.65 over 2017–2021) so that its population base matches the 18–60 sample. After this adjustment, the per-capita consumption of the matched sample tracks the national benchmark closely throughout 2017–2021, indicating that the sample is highly representative of the consumption level of the working-age population nationwide. The dashed vertical line marks the onset of the COVID-19 shock (2020Q1). The series diverge in 2017, when digital payment was still spreading rapidly in China and platform-recorded spending understates total consumption. The treatment and control lines are nearly indistinguishable before 2020 by construction of the matching. Sources: matched Alipay/Huabei panel; NBS quarterly income and consumption releases; NBS annual population communiqués.

Figure 2: Covariate Balance in the Matched Sample




Note: The figure plots the absolute standardized mean difference $|SMD|$ between Huabei users (treatment) and matched non-users (control) for twelve individual characteristics *not* used in the matching procedure, grouped into four semantic dimensions. The vertical dotted line marks $|SMD| = 0.05$; the dashed lines mark the conventional balance thresholds of $|SMD| = 0.10$ and $|SMD| = 0.25$ (Imbens and Rubin, 2015). Log consumption variables are computed from monthly transaction records aggregated to the calendar year and apply the transformation $\ln(\max(c, 1))$. For multi-category covariates (zodiac sign, 12 levels; education, 4 levels), we report the maximum $|SMD|$ across all level indicators. Native, real-estate, and car-ownership flags, as well as the zodiac sign, are observed only for the matched sample; their $|SMD|$ values reflect post-match comparisons of the 217,141 treated and 217,062 matched control consumers.

Figure 3: PPML Event-study (Dynamic DID) Estimates




Note: The figure reports monthly PPML event-study estimates of the treatment effect on total consumption ($\exp(\beta) - 1$, in percent) on the full matched sample of 434,398 consumers, from a single regression with treatment-month interactions over the estimation window December 2018 through December 2021, normalized to December 2019. The window matches the baseline estimation window defined in Section 1.2. The vertical dashed line marks the onset of the COVID-19 shock (January 2020). Shaded areas denote 95% confidence intervals with standard errors clustered at the region level (66 regions). Appendix D.2 reports formal sensitivity of these estimates to pre-trend violations.

Figure 4: Treatment Effect by Liquidity Proxy




Note: PPML treatment-effect estimates (semi-elasticities, expressed in percent) on the full matched sample of 434,398 consumers. Panel (a) sorts consumers into ten deciles by pre-shock (2018–2019) total consumption, from the lowest (1) to the highest (10). Panel (b) sorts consumers into five groups by region of residence, ranked by population-weighted GDP per capita, from the least developed (1) to the most developed (5). Within each cell we estimate a separate PPML regression $Y_{it} = \exp(\beta(\text{Treat}_i \times \text{Post}_t) + \lambda_i + \delta_{ct}) \epsilon_{it}$ with individual and city-by-month fixed effects. Orange triangles report the three-month effect (the acute lockdown window, February–April 2020); blue circles report the two-year effect (2020–2021). Bars are 95% confidence intervals; standard errors are clustered at the region level (66 regions).

Figure 5: Treatment Effect by Payment-Behavior Constraint Proxies




Note: Two-year PPML treatment-effect estimates ($\exp(\beta) - 1$, in percent) with 95% confidence intervals, on the matched sample of 434,398 consumers, with individual and city-by-month fixed effects; standard errors are clustered at the region level (66 regions). The dashed line marks the full-sample two-year effect (8.6%). Panel (a) splits consumers by the number of distinct non-Huabei payment instruments used in 2019 (Alipay balance, debit card, credit card, money-market wallet; an instrument counts if annual spending through it exceeds 10 yuan). Panel (b) splits by active credit-card use in 2019 (at least 500 yuan of card spending). Panel (c) estimates a separate PPML regression for each platform-entry cohort, defined by the year of first active digital-payment use (an exact matching variable, so each cohort is 50/50 treated and control); consumers active before 2010 are pooled, and the 2019 cohort uses a pre-window restricted to 2019H2.

Figure 6: Consumption Response by Category Group and Period




Note: PPML treatment effects ($\exp(\beta) - 1$, in percent) with 95% confidence bands, estimated separately for four non-overlapping post periods (February–April 2020; 2020 Q2; 2020 H2; calendar 2021) against the common pre-window, on the matched sample with individual and city-by-month fixed effects; standard errors are clustered at the region level. In-person merchant categories are assigned to groups by the adjustability of the expenditure; Table 7 lists the composition. The shaded band marks the lockdown quarter, in which aggregate spending contracted and coefficients measure relative declines. Panel (a): consumption categories. Panel (b): committed payments (placebo) and utility-type payment categories.

Figure 7: Credit-Line Usage Among Treated Consumers, by Payment-Instrument Count



Note: Means over the two post years (2020–2021) among treated consumers, by the number of distinct payment instruments used in 2019, with 95% confidence intervals. Panel (a): share of the credit line in use. Panel (b): share of user-months with a positive balance. Panel (c): average outstanding balance (yuan). Panel (d): incidence of delinquency. The annotation in Panel (a) reports each group’s two-year consumption treatment effect for comparison. Balance data are observed for treated consumers only, so the evidence is descriptive; its value is the ordering across groups.

Figure 8: Revolving Balance of Fintech Credit Users



Note: This figure examines the revolving balance among treatment-group consumers (217,141 matched users). The revolving balance is defined as the difference between the total credit line and the effective credit amount. Panel (a) plots the share of fintech credit users carrying a positive revolving balance each month (July 2019–December 2021). Panel (b) shows the average revolving balance (left axis, solid line) and average credit utilization ratio (right axis, dashed line), both conditional on having a positive balance. Panel (c) traces the total revolving balance (left axis, solid line) and number of consumers (right axis, bars) from each consumer’s first month of positive revolving balance, tracking until the balance returns to zero. The median time to repayment is 4 months; 44.6% of consumers clear their balance within 3 months.

Tables

Table 1: Summary Statistics of the Random Sample

	Treated (Huabei users)					Control (non-users)				
	Mean	P10	Median	P90	SD	Mean	P10	Median	P90	SD
<i>Panel A. Individual characteristics</i>										
Age	31.1	20	30	45	9.2	36.3	23	36	50	10.2
Female (= 1)	0.452	—	—	—	0.498	0.475	—	—	—	0.499
Alipay registration year	2013.9	2010	2014	2017	2.9	2013.8	2009	2014	2018	3.4
Taobao registration year	2014.1	2010	2015	2017	2.9	2014.0	2009	2015	2018	3.4
<i>Panel B. Monthly consumption, 2019 (yuan)</i>										
Total consumption	2,655	543	1,859	5,856	2,512	2,345	235	1,415	5,770	2,646
In-person	1,923	296	1,236	4,415	2,107	1,800	117	963	4,607	2,305
Online	732	40	367	1,775	1,101	545	1	186	1,447	1,003
<i>Panel C. Composition within group (% of group)</i>										
Education	Treated		Control		City tier		Treated		Control	
High school or below	78.4		86.0		First-tier		12.4		18.7	
Post-secondary	6.5		2.8		New first-tier		24.0		28.7	
College	13.2		9.4		Second-tier		20.1		21.3	
Graduate (Master's+)	1.9		1.9		Third-tier or below		43.5		31.3	

Note: Statistics for the original random sample of 3,500,000 consumers (499,835 Huabei users and 2,997,291 non-users), before matching. Panel A reports individual characteristics; for the binary female indicator only the group mean (share) is reported. Panel B reports per-consumer average monthly spending in 2019, the pre-shock year. Panel C reports the education and city-tier composition as shares within each group (each column sums to 100%); city tiers follow the standard classification of Chinese prefecture-level cities (first-tier: Beijing, Shanghai, Guangzhou, Shenzhen; new first-tier: 15 cities; second-tier: 30 cities; third-tier or below: all others). The fintech-credit profile of treated users and the full distribution for the matched analysis sample are reported in Appendix Table C1.

Table 2: Covariate Balance: Random Sample vs. Matched Sample (434,398 Consumers)

	Before matching (T=499,835; C=2,997,291)			Matched sample (217,199 pairs) (T=217,199; C=217,199)		
	Treated	Control	$ \Delta_{ND} $	Treated	Control	$ \Delta_{ND} $
<i>Panel A. Matching covariates (9 variables)</i>						
Age	31.1 (9.2)	36.3 (10.2)	0.532	30.4 (8.1)	30.5 (8.0)	0.011
Female (= 1)	0.452	0.475	0.048	0.520	0.520	<0.001
Region (max across 70 dummies)	—	—	0.194	—	—	<0.001
Digital adoption (quarter)	1.30 (1.16)	1.66 (1.85)	0.233	1.17 (0.95)	1.14 (0.85)	0.040
Consumption Dec 2018 (yuan)	2,555 (4,851)	2,054 (4,692)	0.105	2,773 (3,613)	2,723 (3,465)	0.014
Consumption 2019 Q1 (yuan)	6,657 (9,833)	5,774 (10,118)	0.089	7,560 (8,161)	7,467 (7,792)	0.012
Consumption 2019 Q2 (yuan)	7,450 (10,247)	6,671 (10,648)	0.074	8,207 (8,426)	8,138 (8,186)	0.008
Consumption 2019 Q3 (yuan)	8,204 (9,050)	7,441 (9,499)	0.082	8,650 (8,001)	8,602 (7,902)	0.006
Consumption 2019 Q4 (yuan)	9,554 (10,201)	8,256 (10,359)	0.126	10,111 (9,248)	10,046 (9,130)	0.007
<i>Panel B. Held-out covariates (11 variables, not used in matching)</i>						
Consumption 2017 (yuan)	9,686 (18,239)	8,345 (17,705)	0.075	14,043 (20,820)	14,042 (20,081)	<0.001
Consumption 2018 (yuan)	20,834 (29,102)	18,119 (29,617)	0.092	26,095 (28,954)	26,846 (28,591)	0.026
Alipay registration year	2013.92 (2.87)	2013.79 (3.37)	0.041	2013.22 (2.91)	2012.93 (2.99)	0.101
Taobao registration year	2014.05 (2.91)	2013.95 (3.42)	0.031	2013.34 (2.97)	2013.06 (3.06)	0.094
Education						
High school or below	0.784	0.860	0.197	0.736	0.699	0.081
Post-secondary	0.065	0.028	0.179	0.072	0.058	0.058
College	0.132	0.094	0.120	0.164	0.204	0.104
Graduate (Master's or above)	0.019	0.019	0.002	0.028	0.039	0.058
Native to residence city	—	—	—	0.448	0.491	0.087
Real-estate ownership	—	—	—	0.673	0.649	0.050
Car ownership	—	—	—	0.416	0.421	0.011
Zodiac sign (max across 12)	—	—	—	—	—	0.013

Note: Means and absolute normalized differences $|\Delta_{ND}| = |\bar{X}_T - \bar{X}_C| / \sqrt{(S_T^2 + S_C^2)/2}$ for the 3.5M random sample (before matching) and the matched sample of 217,199 pairs (after matching). Standard deviations in parentheses for continuous variables. “< 0.001” indicates exact-matched variables. Dashes (—) indicate variables for which the corresponding statistic is unavailable.

Table 3: Representativeness: Consumption Data Compared to External Benchmark

Year	Sample	National	Urban	Rural	Ratio(%) $\frac{(1) \cdot adj.}{(2)}$	Ratio(%) $\frac{(1) \cdot adj.}{(3)}$	Source
	(1)	(2)	(3)	(4)	(5)	(6)	
2017	15,672	18,322	24,445	10,955	56.3	42.2	NBS
2018	29,076	19,853	26,112	12,124	95.6	72.7	NBS
		15,882	20,890	9,699	90.6	90.8	CFPS
2019	37,560	21,559	28,063	13,328	111.5	85.7	NBS
2020	42,720	21,210	27,007	13,713	125.9	98.9	NBS
2021	48,084	24,100	30,307	15,916	124.7	99.2	NBS

Note: This table shows a comparison of the annual consumption amounts of two groups of consumers in our matched sample, along with data from other sources. We use the annual per capita consumption expenditure from National Bureau of Statistics (NBS), and for 2018, we reference survey data from the China Family Panel Studies (CFPS). Since our study excludes individuals under age 18 and over 60, we adjusted the ratios in columns 5 and 6 according to the population proportion of those aged 18 to 59 for each year. The adjustment ratios for the five years are as follows: 65.9%, 65.3%, 65.3%, 64.0%, and 62.5%.

Table 4: Treatment Effect on Consumption: Levels and Margin Decomposition

	(1) 3 Months 2020/02–04	(2) 6 Months 2020H1	(3) 1 Year Year 2020	(4) 2 Years 2020–2021
<i>Panel A. Total consumption (PPML)</i>				
Treat × Post	0.086*** [9.0%] (0.009)	0.075*** [7.8%] (0.008)	0.081*** [8.5%] (0.009)	0.082*** [8.6%] (0.010)
<i>Panel B. Margin decomposition of total consumption (descriptive)</i>				
Extensive (LPM)	0.018*** (0.001)	0.013*** (0.001)	0.012*** (0.001)	0.014*** (0.002)
Intensive (log)	0.123*** (0.008)	0.121*** (0.007)	0.132*** (0.009)	0.128*** (0.010)
<i>Panel C. By spending channel (PPML)</i>				
In-person	0.140*** [15.0%] (0.010)	0.116*** [12.3%] (0.009)	0.119*** [12.7%] (0.012)	0.122*** [12.9%] (0.013)
Online	−0.021* [−2.1%] (0.012)	−0.009 [−0.9%] (0.012)	−0.001 [−0.1%] (0.012)	−0.003 [−0.3%] (0.012)
<i>Panel D. Decomposition of the in-person effect by frequency and size (descriptive)</i>				
Number of in-person payments (PPML)	0.078*** [8.1%] (0.007)	0.082*** [8.5%] (0.007)	0.068*** [7.1%] (0.007)	0.052*** [5.3%] (0.008)
Amount per in-person payment (residual)	0.062*** [6.4%] (0.006)	0.034*** [3.5%] (0.005)	0.051*** [5.3%] (0.007)	0.070*** [7.2%] (0.008)
Individual FE	Yes	Yes	Yes	Yes
City × Month FE	Yes	Yes	Yes	Yes
SE Clustered	Region	Region	Region	Region
Observations	6,947,248	8,249,857	10,855,075	16,065,511

Note: Treatment effect of fintech credit access on consumption for 434,398 matched consumers. Each cell reports the coefficient on Treat×Post; for the PPML panels (A, C, and D) the bracketed value is the implied semi-elasticity $(\exp \beta - 1) \times 100$. Panel A is the causally identified total effect; Panel B is a descriptive decomposition, since the positive-consumption subsample in the intensive margin is itself affected by treatment; the intensive-margin row is estimated on the positive-consumption subsample. Panel D decomposes the Panel C in-person effect into transaction frequency and average size: the number-of-in-person-payments coefficient is the PPML estimate on in-person transaction counts, and amount per in-person payment is the residual. The standard error of the residual is obtained by a region block bootstrap ($B = 200$, resampling the 66 regions with replacement) and accounts for the covariance between the two PPML estimates. All specifications include individual and city×month fixed effects; standard errors clustered at the region level (66 regions) in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 5: External Validity: PPML Treatment Effect Across Estimation Specifications and Post-Windows

Specification	N individuals	PPML semi-elasticity ($\exp(\hat{\beta}) - 1$)			
		3 months	6 months	Year 2020	2 years
Baseline (matched, panel FE)	434,398	9.03%*** (0.0093)	7.76%*** (0.0080)	8.46%*** (0.0092)	8.56%*** (0.0103)
Specification 1: IPW (collapsed PPML)	2,987,387	14.91%*** (0.0106)	13.03%*** (0.0099)	14.34%*** (0.0120)	14.75%*** (0.0137)
Specification 2: PSM 1:1 ($c = 0.10$ SD)	902,618	19.98%*** (0.0092)	18.14%*** (0.0081)	19.99%*** (0.0100)	20.93%*** (0.0114)
Specification 2: PSM 1:3 ($c = 0.25$ SD)	1,459,240	19.75%*** (0.0084)	18.20%*** (0.0077)	20.13%*** (0.0095)	21.06%*** (0.0112)
Specification 2: PSM 1:5 ($c = 0.50$ SD)	1,826,294	19.79%*** (0.0083)	18.18%*** (0.0077)	20.09%*** (0.0096)	20.99%*** (0.0112)
Specification 3: reweighted matched	434,398	10.05%*** (0.0114)	8.24%*** (0.0094)	8.95%*** (0.0104)	8.60%*** (0.0111)

Notes: All estimates use pre-period months 24–36 (December 2018 plus all of 2019). Post-windows: 3 months (Feb–Apr 2020), 6 months (2020H1), Year 2020, and 2 years (2020–2021). Baseline and Specification 3 use `fepois` with individual and city×month fixed effects on the matched panel; Specification 2 uses `fepois` with individual and month fixed effects on PSM-matched panels (matching weights $1/K$ applied to controls); Specification 1 uses collapsed quasi-Poisson with pre-period consumption control on the IPW-weighted 3.5M sample (panel `fepois` on 3M individuals is memory-prohibitive). Region-clustered standard errors (66 regions) in parentheses. *** $p < 0.01$. Companion LPM estimates are reported in Appendix D.1.

Table 6: Heterogeneous Effects by Liquidity Constraints

	Coefficient	(SE)	Implied effect (%)	Users
<i>Panel A. Payment instruments used in 2019</i>				
1 instrument	0.242***	(0.017)	27.3	12,142
2 instruments	0.152***	(0.013)	16.4	43,482
3 instruments	0.109***	(0.012)	11.5	161,299
4 instruments	0.057***	(0.010)	5.9	217,037
<i>Panel B. Credit-card activity in 2019 (cutoff 500 yuan/yr)</i>				
No active use	0.175***	(0.011)	19.2	103,711
Active use	0.060***	(0.010)	6.2	330,540
<i>Panel C. Platform-entry cohort (year of first digital payment)</i>				
Pre-2010	0.030***	(0.009)	3.1	46,000
2010	0.058***	(0.014)	6.0	22,099
2011	0.062***	(0.008)	6.4	30,310
2012	0.072***	(0.013)	7.5	32,226
2013	0.073***	(0.012)	7.6	36,010
2014	0.086***	(0.016)	9.0	39,933
2015	0.074***	(0.009)	7.7	36,218
2016	0.081***	(0.017)	8.5	22,960
2017	0.126***	(0.015)	13.4	145,172
2018	0.132***	(0.015)	14.1	16,223
2019	0.159***	(0.025)	17.2	7,052

Note: Two-year PPML coefficients on Treat×Post with region-clustered standard errors in parentheses; implied effects are $(\exp \beta - 1) \times 100$. Panel A splits the matched sample by the number of distinct non-Huabei payment instruments used in 2019; Panel B by active credit-card use in 2019 (at least 500 yuan of card spending); Panel C estimates separate regressions by platform-entry cohort (year of first digital-payment activity; consumers active before 2010 are pooled, and the 2019 cohort uses a pre-window restricted to 2019H2). All specifications include individual and city×month fixed effects. The income and regional gradients are reported in Figure 4. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 7: Consumption Response by Category Group and Period

	Non-overlapping periods				Two years	2yr effect
	Feb–Apr 20	2020 Q2	2020 H2	2021	2020–21	(%)
<i>Panel A. PPML treatment effects by category group and period</i>						
Daily necessities	0.074*** (0.016)	0.063*** (0.015)	0.133*** (0.016)	0.140*** (0.021)	0.119*** (0.017)	12.6
Health & medical	0.056*** (0.018)	0.053** (0.021)	0.068*** (0.016)	0.045*** (0.016)	0.051*** (0.013)	5.3
Discretionary	0.238*** (0.016)	0.122*** (0.013)	0.143*** (0.014)	0.132*** (0.014)	0.137*** (0.012)	14.7
Big-ticket & travel	0.112*** (0.028)	0.001 (0.020)	-0.003 (0.018)	-0.017 (0.021)	-0.002 (0.018)	-0.2
Committed (placebo)	0.057 (0.047)	0.025 (0.042)	0.086*** (0.020)	0.064*** (0.023)	0.061*** (0.020)	6.3
Utility-type	0.139*** (0.028)	0.124*** (0.019)	0.124*** (0.020)	0.121*** (0.026)	0.121*** (0.021)	12.8
<i>Panel B. Group composition</i>						
	MCC codes		Share of offline spending (%)			
Daily necessities	3		10.6			
Health & medical	5		4.1			
Discretionary	10		27.2			
Big-ticket & travel	12		7.2			
Committed (placebo)	4		1.3			
Utility-type	5		2.4			
Named categories, total	39		53.0			
Unclassified small categories	–		47.0			

Note: Panel A reports PPML coefficients on Treat×Post with region-clustered standard errors in parentheses, estimated separately for each post period against the common pre-window (December 2018 through December 2019); the final column converts the two-year coefficient to a percentage effect. In-person merchant-category codes are assigned to groups by the adjustability of the underlying expenditure; outlier categories (travel agencies, mahjong parlors, health supplements) are excluded. Panel B reports the number of merchant-category codes in each group and the group's share of total in-person spending; named categories cover 53% of in-person spending. All specifications include individual and city×month fixed effects. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 8: Consumption versus Payment-Channel Migration: Margin Decomposition

	Pre-active share (%)	Intensive margin (PPML)		Extensive margin (LPM, pp)	
		Coefficient	Effect (%)	Full sample	Dormant users
Daily necessities	98	0.117*** (0.017)	12.4	1.65*** (0.22)	4.79*** (0.52)
Health & medical	80	-0.017 (0.014)	-1.7	0.37*** (0.12)	1.55*** (0.14)
Discretionary	100	0.137*** (0.012)	14.6	3.50*** (0.20)	4.76*** (1.02)
Big-ticket & travel	83	-0.078*** (0.019)	-7.5	-0.06 (0.11)	1.99*** (0.11)
Committed (placebo)	46	-0.275*** (0.042)	-24.1	0.36*** (0.09)	0.14 (0.10)
Utility-type	70	0.014 (0.016)	1.4	0.79*** (0.16)	0.95*** (0.08)

Note: The intensive margin re-estimates the two-year PPML effect on the subsample of consumers with positive pre-window spending in the category group; for these consumers payment-channel substitution cannot generate measured growth. The extensive margin reports linear-probability estimates (in percentage points) of the treatment effect on the incidence of any spending in the group, for the full sample and for consumers with zero pre-window spending (“dormant”). The first column reports the share of consumers with positive pre-window spending. Standard errors clustered at the region level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 9: Alternative Explanations and Evidence

Alternative explanation	Test	Result
Pure borrowing channel	Credit-line usage among treated users (Figure 7)	86% of treated users never carry a positive installment balance (93% among single-instrument users); the most constrained users borrow the least yet respond the most
Channel migration	Intensive/extensive decomposition (Table 8)	Effects survive among pre-active users where migration is impossible; utility-type categories show no intensive-margin effect
Income or wealth effect	Category composition (Table 7)	Response tilts toward daily necessities and constrained users, inconsistent with a proportional wealth effect; committed payments do not respond
Promotion targeting	Baseline effect controlling for pre-period promotion exposure	Overall effect survives (coef. 0.072, SE 0.009)
Temporary pandemic relief	Effects in 2021 only, after pandemic disruption ended	Effect persists (coef. 0.094, SE 0.013); necessity effect is largest in 2021

Note: Summary of alternative explanations for the consumption response and the evidence that bounds each. The promotion-control estimate is a PPML regression adding an interaction of treatment with pre-period promotion exposure; the persistence estimate restricts the post window to calendar 2021.

References

- Abadie, Alberto and Guido W. Imbens (2006). “Large Sample Properties of Matching Estimators for Average Treatment Effects”. en. In: *Econometrica* 74.1, pp. 235–267.
- Acharya, Viral V., Heitor Almeida, Filippo Ippolito, and Ander Perez (2014). “Credit lines as monitored liquidity insurance: Theory and evidence”. In: *Journal of Financial Economics* 112.3, pp. 287–319.
- Agarwal, Sumit, Shashwat Alok, Pulak Ghosh, and Sudip Gupta (2019). “Financial Inclusion and Alternate Credit Scoring: Role of Big Data and Machine Learning in Fintech”. en. In: *Journal of Money, Credit and Banking*. Forthcoming.
- Agarwal, Sumit, Shashwat Alok, Pulak Ghosh, and Xiaoyu Zhang (2024). *Why do financially unconstrained individuals respond to higher credit limits?* Working paper 4954242. SSRN.
- Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel (2018). “Do Banks Pass through Credit Expansions to Consumers Who Want to Borrow?” In: *The Quarterly Journal of Economics* 133.1, pp. 129–190.
- Agarwal, Sumit, Pulak Ghosh, Jing Li, and Tianyue Ruan (2024). “Digital Payments and Consumption: Evidence from the 2016 Demonetization in India”. en. In: *The Review of Financial Studies* 37.8, pp. 2550–2585.
- Agarwal, Sumit and Wenlan Qian (2014). “Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore”. en. In: *American Economic Review* 104.12, pp. 4205–4230.
- Aydin, Deniz (2022). “Consumption Response to Credit Expansions: Evidence from Experimental Assignment of 45,307 Credit Lines”. en. In: *American Economic Review* 112.1, pp. 1–40.
- Bachas, Pierre, Paul Gertler, Sean Higgins, and Enrique Seira (2021). “How Debit Cards Enable the Poor to Save More”. In: *The Journal of Finance* 76.4, pp. 1913–1957.
- Baker, Scott R. (2018). “Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data”. en. In: *Journal of Political Economy* 126.4, pp. 1504–1557.
- Baker, Scott R., Robert A. Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis (2020). “How Does Household Spending Respond to an Epidemic? Consumption during the 2020 COVID-19 Pandemic”. In: *The Review of Asset Pricing Studies* 10.4, pp. 834–862.
- Baker, Scott R. and Constantine Yannelis (2017). “Income changes and consumption: Evidence from the 2013 federal government shutdown”. en. In: *Review of Economic Dynamics* 23, pp. 99–124.
- Baugh, Brian, Itzhak Ben-David, Hoonsuk Park, and Jonathan A. Parker (2021). “Asymmetric Consumption Smoothing”. en. In: *American Economic Review* 111.1, pp. 192–230.

- Berg, Tobias, Valentin Burg, Ana Gombović, and Manju Puri (2020). “On the Rise of FinTechs: Credit Scoring Using Digital Footprints”. en. In: *The Review of Financial Studies* 33.7, pp. 2845–2897.
- Berg, Tobias, Valentin Burg, Jan Keil, and Manju Puri (2025). “The Economics of “Buy Now, Pay Later”: A Merchant’s Perspective”. en. In: *Journal of Financial Economics* 171, p. 104093.
- Berg, Tobias, Andreas Fuster, and Manju Puri (2022). “FinTech Lending”. en. In: *Annual Review of Financial Economics* 14.1, pp. 187–207.
- Bharadwaj, Prashant, William Jack, and Tavneet Suri (2021). “Fintech and household resilience to shocks: Evidence from digital loans in Kenya”. In: *Journal of Development Economics* 153, p. 102697.
- Bian, Wenlong, Lin William Cong, and Yang Ji (2024). *The Rise of E-Wallet Super-Apps and Buy-Now-Pay-Later*. en. Tech. rep. w33178. Cambridge, MA: National Bureau of Economic Research.
- Braxton, J. Carter, Kyle F. Herkenhoff, and Gordon M. Phillips (2024). “Can the Unemployed Borrow? Implications for Public Insurance”. In: *Journal of Political Economy* 132.9, pp. 3025–3076.
- 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.
- Carroll, Christopher D (1997). “Buffer-stock saving and the life cycle/permanent income hypothesis”. In: *Quarterly Journal of Economics* 112.1, pp. 1–55.
- Chen, Haiqiang, Wenlan Qian, and Qiang Wen (2021). “The Impact of the COVID-19 Pandemic on Consumption: Learning from High-Frequency Transaction Data”. en. In: *AEA Papers and Proceedings* 111, pp. 307–311.
- Chen, Jiafeng and Jonathan Roth (2024). “Logs with Zeros? Some Problems and Solutions”. en. In: *The Quarterly Journal of Economics* 139.2, pp. 891–936.
- Chen, Tao, Yi Huang, Chen Lin, and Zixia Sheng (2022). “Finance and Firm Volatility: Evidence from Small Business Lending in China”. en. In: *Management Science* 68.3, pp. 2226–2249.
- Chetty, Raj, John N Friedman, Michael Stepner, and the Opportunity Insights Team (2024). “The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data”. In: *The Quarterly Journal of Economics* 139.2, pp. 829–889.
- Cox, Natalie, Peter Ganong, Pascal Noel, Joseph Vavra, Arlene Wong, Diana Farrell, Fiona Greig, and Erica Deadman (2020). “Initial Impacts of the Pandemic on Consumer Behavior: Evidence from Linked Income, Spending, and Savings Data”. In: *Brookings Papers on Economic Activity* 51.2, pp. 35–82.
- Crouzet, Nicolas, Apoorv Gupta, and Filippo Mezzanotti (2023). “Shocks and Technology Adoption: Evidence from Electronic Payment Systems”. In: *Journal of Political Economy* 131.11, pp. 3003–3065.

- Crump, Richard K, V Joseph Hotz, Guido W Imbens, and Oscar A Mitnik (2009). “Dealing with limited overlap in estimation of average treatment effects”. In: *Biometrika* 96.1, pp. 187–199.
- Deaton, Angus (1991). “Saving and Liquidity Constraints”. In: *Econometrica* 59.5, pp. 1221–1248.
- deHaan, Ed, Jungbae Kim, Ben Lourie, and Chenqi Zhu (2024). “Buy Now Pay (Pain?) Later”. en. In: *Management Science* 70.8, pp. 5586–5598.
- Di Maggio, Marco and Dimuthu Ratnadiwakara (2026). “Invisible Primes: Fintech Lending with Alternative Data”. en. In: *Management Science*. Articles in Advance.
- Di Maggio, Marco, Emily Williams, and Justin Katz (2022). *Buy Now, Pay Later Credit: User Characteristics and Effects on Spending Patterns*. en. Tech. rep. w30508. Cambridge, MA: National Bureau of Economic Research, w30508.
- Di Maggio, Marco and Vincent Yao (2021). “Fintech Borrowers: Lax Screening or Cream-Skimming?” en. In: *The Review of Financial Studies* 34.10, pp. 4565–4618.
- Didier, Tatiana, Erik H.B. Feyen, Ruth Llovet Montanes, and Oya Pinar Ardic Alper (2022). “Global patterns of fintech activity and enabling factors”. In: *World Bank Group Fintech and the Future of Finance report*.
- Ding, Jing, Lei Jiang, Lucy Msall, and Matthew J. Notowidigdo (2025). “Consumer-Financed Fiscal Stimulus: Evidence from Digital Coupons in China”. In: *American Economic Review: Insights* 7.3, pp. 411–427.
- Dong, Yanfei, Jiayin Hu, Yiping Huang, Han Qiu, and Yingguang Zhang (2025). *The Use and Disuse of FinTech Credit: When Buy-Now-Pay-Later Meets Credit Reporting*. en. BIS Working Papers 1239. Bank for International Settlements.
- Eça, Afonso, Miguel Almeida Ferreira, Melissa Porras Prado, and Antonino Emanuele Rizzo (2022). *The Real Effects of FinTech Lending on SMEs: Evidence from Loan Applications*. en. Working paper 2639. European Central Bank.
- Frost, Jon, Leonardo Gambacorta, Yi Huang, Hyun Song Shin, and Pablo Zbinden (2019). “BigTech and the changing structure of financial intermediation”. en. In: *Economic Policy* 34.100, pp. 761–799.
- Fulford, Scott L. and Scott D. Schuh (2024). “Credit cards, credit utilization, and consumption”. en. In: *Journal of Monetary Economics*, p. 103619.
- 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.
- Ganong, Peter and Pascal Noel (2019). “Consumer Spending during Unemployment: Positive and Normative Implications”. en. In: *American Economic Review* 109.7, pp. 2383–2424.
- Gelman, Michael, Shachar Kariv, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis (2020). “How individuals respond to a liquidity shock: Evidence from the 2013 government shut-down”. en. In: *Journal of Public Economics* 189, p. 103917.

- Gelman, Michael and Nikolai Roussanov (2024). “Managing Mental Accounts: Payment Cards and Consumption Expenditures”. en. In: *The Review of Financial Studies* 37.8, pp. 2586–2624.
- Gertler, Paul, Brett Green, and Catherine Wolfram (2024). “Digital Collateral”. In: *The Quarterly Journal of Economics* 139.3, pp. 1713–1766.
- Ghosh, Pulak, Boris Vallée, and Yao Zeng (2026). “FinTech Lending and Cashless Payments”. In: *The Journal of Finance* 81.2, pp. 1053–1101.
- Gross, David B. and Nicholas S. Souleles (2002). “Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data”. en. In: *The Quarterly Journal of Economics* 117.1, pp. 149–185.
- Gruber, Jonathan (1997). “The Consumption Smoothing Benefits of Unemployment Insurance”. In: *The American Economic Review* 87.1, p. 192.
- Hartman, Erin, Richard Grieve, Roland Ramsahai, and Jasjeet S Sekhon (2015). “From sample average treatment effect to population average treatment effect on the treated”. In: *Journal of the Royal Statistical Society: Series A* 178.3, pp. 757–778.
- Hau, Harald, Yi Huang, Chen Lin, Hongzhe Shan, Zixia Sheng, and Lai Wei (2024). “FinTech Credit and Entrepreneurial Growth”. en. In: *The Journal of Finance* 79.5, pp. 3309–3359.
- Herkenhoff, Kyle F. (2019). “The Impact of Consumer Credit Access on Unemployment”. In: *The Review of Economic Studies* 86.6, pp. 2605–2642.
- Higgins, Sean (2024). “Financial Technology Adoption: Network Externalities of Cashless Payments in Mexico”. In: *American Economic Review* 114.11, pp. 3469–3512.
- Holmström, Bengt and Jean Tirole (1998). “Private and Public Supply of Liquidity”. In: *Journal of Political Economy* 106.1, pp. 1–40.
- Hu, Jiayin, Yiping Huang, and Jialin Liu (2024). “The changing face of consumer credit: Evidence from a big tech platform in China”. In: *Pacific-Basin Finance Journal* 83, p. 102254.
- Hundtofte, Sean, Arna Olafsson, and Michaela Pagel (2019). *Credit Smoothing*. en. Tech. rep. w26354. Cambridge, MA: National Bureau of Economic Research, w26354.
- Imbens, Guido W and Jeffrey M Wooldridge (2009). “Recent developments in the econometrics of program evaluation”. In: *Journal of Economic Literature* 47.1, pp. 5–86.
- Imbens, Guido W. (2015). “Matching Methods in Practice: Three Examples”. en. In: *Journal of Human Resources* 50.2, pp. 373–419.
- Imbens, Guido W. and Donald B. Rubin (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. 1st ed. Cambridge University Press.
- Jack, William and Tavneet Suri (2014). “Risk Sharing and Transactions Costs: Evidence from Kenya’s Mobile Money Revolution”. en. In: *American Economic Review* 104.1, pp. 183–223.
- Jaffe, Sonia, Anup Malani, and Julian Reif (2026). “Access to credit reduces the value of insurance”. In: *Journal of Public Economics* 258, p. 105633.


- Jappelli, Tullio and Luigi Pistaferri (2010). “The Consumption Response to Income Changes”. en. In: *Annual Review of Economics* 2.1, pp. 479–506.
- Ji, Yang, Xue Wang, Yiping Huang, Shu Chen, and Fang Wang (2023). “Buy now, pay later as liquidity insurance: Evidence from an early experiment in China”. In: *China Economic Review* 80, p. 101998.
- Kaplan, Greg and Giovanni L. Violante (2014). “A Model of the Consumption Response to Fiscal Stimulus Payments”. In: *Econometrica* 82.4, pp. 1199–1239.
- Kim, Olivia S., Jonathan A. Parker, and Antoinette Schoar (2025). “Revenue Collapses and the Consumption of Small Business Owners in the COVID-19 Pandemic”. en. In: *Journal of Financial Economics* 170, p. 104079.
- Lee, Jung Youn, Joonhyuk Yang, and Eric T. Anderson (2025). “Using Grocery Data for Credit Decisions”. en. In: *Management Science* 71.4, pp. 2753–2777.
- Li, Jian and Stefano Pegoraro (2026). “Borrowing from a Bigtech Platform”. In: *The Review of Financial Studies*. Advance article.
- Liu, Lei, Guangli Lu, and Wei Xiong (2022). *The Big Tech Lending Model*. en. Tech. rep. w30160. Cambridge, MA: National Bureau of Economic Research, w30160.
- Liu, Qiao, Qiaowei Shen, Zhenghua Li, and Shu Chen (2021). “Stimulating Consumption at Low Budget: Evidence from a Large-Scale Policy Experiment amid the COVID-19 Pandemic”. In: *Management Science* 67.12, pp. 7291–7307.
- Lunt, Mark (2014). “Selecting an Appropriate Caliper Can Be Essential for Achieving Good Balance With Propensity Score Matching”. en. In: *American Journal of Epidemiology* 179.2, pp. 226–235.
- Lusardi, Annamaria (1996). “Permanent Income, Current Income, and Consumption: Evidence From Two Panel Data Sets”. en. In: *Journal of Business & Economic Statistics* 14.1, pp. 81–90.
- Ouyang, Shumiao (2026). “Cashless Payment and Financial Inclusion”. en. In: *Journal of Financial Economics* 180, p. 104277.
- Parker, Jonathan A., Jake Schild, Laura Erhard, and David S Johnson (2022). “Economic impact payments and household spending during the pandemic”. In: *Brookings Papers on Economic Activity* 2022.2, pp. 81–156.
- Rambachan, Ashesh and Jonathan Roth (2023). “A More Credible Approach to Parallel Trends”. In: *The Review of Economic Studies* 90.5, pp. 2555–2591.
- Rosenzweig, Mark R. and Kenneth I. Wolpin (1993). “Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investments in Bullocks in India”. en. In: *Journal of Political Economy* 101.2, pp. 223–244.
- Silva, J. M. C. Santos and Silvana Tenreiro (2006). “The Log of Gravity”. en. In: *The Review of Economics and Statistics* 88.4, pp. 641–658.

- Stuart, Elizabeth A, Stephen R Cole, Catherine P Bradshaw, and Philip J Leaf (2011). “The use of propensity scores to assess the generalizability of results from randomized trials”. In: *Journal of the Royal Statistical Society: Series A* 174.2, pp. 369–386.
- Sufi, Amir (2009). “Bank Lines of Credit in Corporate Finance: An Empirical Analysis”. In: *The Review of Financial Studies* 22.3, pp. 1057–1088.
- Suri, Tavneet and William Jack (2016). “The long-run poverty and gender impacts of mobile money”. In: *Science* 354.6317, pp. 1288–1292.
- Tang, Huan (2019). “Peer-to-Peer Lenders versus Banks: Substitutes or Complements?” In: *The Review of Financial Studies* 32.5, pp. 1900–1938.
- Townsend, Robert M. (1994). “Risk and Insurance in Village India”. In: *Econometrica* 62.3, p. 539.
- Vallée, Boris and Yao Zeng (2019). “Marketplace Lending: A New Banking Paradigm?” In: *The Review of Financial Studies* 32.5, pp. 1939–1982.
- Vissing-Jorgensen, Annette (2011). *Consumer credit: learning your customer’s default risk from what (s)he buys*. Working paper. Northwestern University.
- Yin, Xiao (2026). “Learning in the Limit: Income Inference from Credit Extensions”. In: *The Journal of Finance*. Forthcoming.
- Zeldes, Stephen P. (1989). “Consumption and Liquidity Constraints: An Empirical Investigation”. In: *Journal of Political Economy* 97.2, pp. 305–346.


A Additional Figures and Tables

This appendix collects supplementary figures and tables referenced in the main text.

Figure A9: Adoption-Adjusted Dynamic Effects: Total and by Channel




(a) Total consumption, with monthly count of new control-group adopters (right axis)



(b) Effect by channel: in-person vs. online (e-commerce)

Note: Monthly event-study PPML estimates on the full matched sample of 434,398 consumers, normalized to December 2019 (the dashed vertical line marks the January 2020 onset of the COVID-19 shock). The blue “Baseline” series fixes treatment status as of the pre-shock period; the orange “Adoption-adjusted” series re-classifies control consumers who first adopt Huabei after the shock as treated from their adoption month forward, following the definition in column (2) of Table A14. The light-grey bars in panel (a) report the number of control consumers who newly adopt Huabei each month (right axis, capped at 8,000). All specifications include individual and city-by-month fixed effects; standard errors are clustered at the region level (66 regions); shaded bands are 95% confidence intervals.

Figure A10: Dynamic Effect on In-Person Consumption, by Category Group



Note: Monthly event-study estimates of the effect of Huabei access on in-person spending, by category group, on the matched sample. Panel (a) groups merchant categories by durability; panel (b) by necessity, with restaurants shown separately. Coefficients are normalized to December 2019; the dashed line marks the January 2020 COVID-19 onset; shaded bands are 95% confidence intervals.

Table A10: Treatment Effect by Consumption Category

Spending	Category	3-Month (2020/02–04)			2-Year (2020–2021)		
		Effect (%)	Baseline (CNY/mo)	Increase (CNY/mo)	Effect (%)	Baseline (CNY/mo)	Increase (CNY/mo)
SND	Food Away From Home	19.51*** (1.46)	187.0	36.5	13.15*** (1.30)	345.1	45.4
ND	Retail Nondurables	18.10*** (1.73)	168.9	30.6	20.99*** (2.24)	290.2	60.9
ND	Other Nondurables	16.84*** (1.80)	158.7	26.7	10.70*** (1.34)	265.2	28.4
SND	Groceries	8.46*** (1.46)	227.5	19.2	14.36*** (1.92)	257.6	37.0
SND	Prof. & Pers. Services	15.83*** (2.72)	79.5	12.6	16.37*** (2.27)	142.1	23.3
D	Entertainment	27.82*** (4.01)	31.2	8.7	17.55*** (1.91)	52.0	9.1
D	Retail Durables	13.51*** (2.25)	62.9	8.5	13.31*** (2.17)	107.0	14.2
D	Auto Repair	23.84*** (5.35)	35.3	8.4	14.75*** (2.44)	55.2	8.1
D	Home Improvement	12.39*** (3.94)	48.5	6.0	−1.66 (2.73)	91.1	−1.5
ND	Department Stores	22.63*** (4.44)	25.4	5.7	19.12*** (2.95)	56.2	10.7
ND	Medical / Healthcare	6.93*** (1.70)	56.6	3.9	7.29*** (1.41)	99.7	7.3
SND	Ground Transportation	9.09*** (1.47)	40.4	3.7	7.80*** (1.04)	93.5	7.3
D	Hotels & Rental Cars	17.64*** (5.44)	18.0	3.2	8.76*** (2.97)	55.0	4.8
D	Insurance	9.15*** (2.41)	25.6	2.3	6.31*** (1.56)	26.7	1.7
SND	Utilities	18.70* (10.54)	4.3	0.8	15.64*** (3.82)	7.2	1.1
D	Other Durables	4.36 (4.75)	12.8	0.6	2.19 (2.93)	51.8	1.1
SND	Flights	17.22 (18.18)	0.5	0.1	11.52 (8.32)	1.3	0.1
ND	Pharmacies	−0.15 (3.11)	23.7	−0.0	5.43** (2.49)	28.9	1.6
SND	Telecom	−1.08 (4.48)	11.5	−0.1	14.10*** (3.96)	17.7	2.5
Total (19 categories)			1218.3	177.3		2043.5	263.2

Notes: PPML treatment effects on offline spending by category, for 434,398 matched consumers. Effect is the PPML semi-elasticity (%); Baseline is the control-group mean monthly spending (CNY) in the post window; Increase = Effect $\times$ Baseline / 100. Standard errors (in parentheses) clustered at the region level. Spending type: ND = non-durable, SND = semi-non-durable, D = durable. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A11: External Credit Capacity

	(1) 2020/02–04	(2) 2020H1	(3) Year 2020	(4) 2020–2021
<i>PPML</i>				
<i>Treat</i> × <i>Post</i> × Low external credit	0.089*** (0.021)	0.069*** (0.020)	0.078*** (0.016)	0.077*** (0.016)
[% effect]	[9.3%]	[7.2%]	[8.1%]	[8.0%]
<i>Treat</i> × <i>Post</i> × High external credit	0.097*** (0.030)	0.041 (0.025)	0.016 (0.020)	−0.009 (0.019)
[% effect]	[10.2%]	[4.2%]	[1.6%]	[−0.9%]
Individual FE	Yes	Yes	Yes	Yes
City × Month FE	Yes	Yes	Yes	Yes
SE Clustered	Region	Region	Region	Region
Num. Obs.	217,744	258,571	340,225	503,533

Table A12: Additional Heterogeneity: Education

	(1)	(2)	(3)	(4)
	High school or below	Post-secondary	College	Graduate
<i>Panel A: PPML (1 Year)</i>				
<i>Treat × Post</i>	0.087***	0.065***	0.081***	0.057***
	(0.011)	(0.012)	(0.008)	(0.012)
[% effect]	[9.1%]	[6.7%]	[8.4%]	[5.8%]
Individual FE	Yes	Yes	Yes	Yes
City × Month FE	Yes	Yes	Yes	Yes
SE Clustered	Region	Region	Region	Region
Num. Obs.	7,788,025	704,750	1,999,275	362,025
Share of Obs.	71.8%	6.5%	18.4%	3.3%
<i>Panel B: PPML (2 Years)</i>				
<i>Treat × Post</i>	0.083***	0.081***	0.087***	0.058***
	(0.012)	(0.011)	(0.011)	(0.014)
[% effect]	[8.7%]	[8.4%]	[9.1%]	[5.9%]
Individual FE	Yes	Yes	Yes	Yes
City × Month FE	Yes	Yes	Yes	Yes
SE Clustered	Region	Region	Region	Region
Num. Obs.	11,526,277	1,043,028	2,958,927	535,797
Share of Obs.	71.8%	6.5%	18.4%	3.3%

Table A13: Additional Heterogeneity: Financial Sophistication Index

A. Financial Sophistication Index as an Interaction Term				
	(1)	(2)	(3)	(4)
	2020/02–04	2020H1	Year 2020	2020–2021
<i>PPML</i>				
<i>Treat</i> × <i>Post</i>	0.125*** (0.012)	0.098*** (0.011)	0.102*** (0.012)	0.099*** (0.013)
<i>Treat</i> × <i>Post</i> × <i>FS</i>	−0.045*** (0.007)	−0.019*** (0.006)	−0.019*** (0.006)	−0.015** (0.006)
Individual FE	Yes	Yes	Yes	Yes
City × Month FE	Yes	Yes	Yes	Yes
SE Clustered	Region	Region	Region	Region
Num. Obs.	10,489,932	11,655,480	13,986,576	18,648,768
B. Low vs. High Financial Sophistication (PPML)				
	3 Months		2 Years	
	(1) Low FS	(2) High FS	(3) Low FS	(4) High FS
<i>Treat</i> × <i>Post</i>	0.112*** (0.011)	0.088*** (0.009)	0.092*** (0.012)	0.087*** (0.010)
[% effect]	[11.8%]	[9.2%]	[9.6%]	[9.1%]
<i>Difference (Triple-Diff): High – Low</i>				
	−0.035*** (0.006)		−0.004 (0.005)	
Individual FE	Yes	Yes	Yes	Yes
City × Month FE	Yes	Yes	Yes	Yes
SE Clustered	Region	Region	Region	Region
Num. Obs.	5,244,902	5,244,939	9,324,278	9,324,336
Share of Obs.	50.0%	50.0%	50.0%	50.0%

Table A14: The Adoption and Exit of Fintech Credit

	(1) Baseline	(2) Adoption	(3) Exit	(4) Adoption & Exit
<i>PPML</i>				
<i>Treat</i> × <i>Post</i>	0.097*** (0.011)	0.128*** (0.012)	0.164*** (0.014)	0.183*** (0.015)
[% effect]	[10.2%]	[13.6%]	[17.8%]	[20.0%]
Individual FE	Yes	Yes	Yes	Yes
City × Month FE	Yes	Yes	Yes	Yes
SE Clustered	Region	Region	Region	Region
Num. Obs.	20,841,648	20,841,648	20,841,648	20,841,648

Table A15: Robustness in the External-Credit-Observed Subsample

	(1)	(2)	(3)	(4)
	2020/02–04	2020H1	Year 2020	2020–2021
<i>Credit Subsample (PPML)</i>				
<i>Treat</i> × <i>Post</i>	0.114***	0.086***	0.082***	0.060***
	(0.022)	(0.023)	(0.020)	(0.020)
[% effect]	[12.0%]	[9.0%]	[8.5%]	[6.1%]
Individual FE	Yes	Yes	Yes	Yes
City × Month FE	Yes	Yes	Yes	Yes
SE Clustered	Region	Region	Region	Region
Num. Obs.	207,893 individuals with observable credit lines			

Note: PPML estimates on the subsample of matched consumers with observable external bank credit lines. This subsample confirms the baseline effect is robust and slightly smaller, consistent with this subsample being less credit-constrained on average. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A16: Heterogeneity Analysis: Age, Gender, and Occupation

Panel A: Gender (PPML, 2 Years)						
	(1) Male	(2) Female				
<i>Treat × Post</i>	0.099*** (0.013)	0.095*** (0.012)				
[% effect]	[10.4%]	[10.0%]				
Num. Obs.	9,999,143	10,842,336				
Panel B: Age Groups (PPML, 2 Years)						
	(1) Under 25	(2) 26–30	(3) 31–35	(4) 36–40	(5) 41–45	(6) 45+
<i>Treat × Post</i>	0.094*** (0.008)	0.096*** (0.013)	0.091*** (0.015)	0.100*** (0.014)	0.105*** (0.018)	0.077*** (0.023)
[% effect]	[9.9%]	[10.1%]	[9.5%]	[10.5%]	[11.0%]	[8.0%]
Panel C: Occupation (PPML, 2 Years)						
	(1) White collar	(2) Blue collar	(3) Students	(4) Teachers	(5) Civil servants	(6) Medical staff
<i>Treat × Post</i>	0.073*** (0.010)	0.163*** (0.012)	0.098*** (0.010)	0.083*** (0.012)	0.052** (0.024)	0.111*** (0.019)
[% effect]	[7.6%]	[17.7%]	[10.3%]	[8.7%]	[5.3%]	[11.7%]
Individual FE				Yes		
City × Month FE				Yes		
SE Clustered				Region (66 regions)		

Note: PPML estimates on the corrected 434K matched sample. Blue collar workers show the largest effect (17.7%), consistent with the income channel. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

B Balance across Granular Merchant Categories

A potential concern with our matching procedure is that, despite balancing on aggregate spending and demographic characteristics, the treatment and control groups may still differ in the *composition* of their spending. We therefore assess balance across 42 merchant category codes (MCCs) spanning six economic dimensions, ranging from inelastic-demand categories to durable goods. Because our sample is large enough that conventional t -tests flag economically trivial differences as statistically significant, we measure balance with the scale-free normalized difference (ND) of [Imbens and Wooldridge \(2009\)](#), and treat $|\Delta_{\text{ND}}| < 0.10$ as the conventional threshold for excellent balance.¹⁷

Table B17 reports the ND for each MCC along three dimensions: the unconditional mean $E(Y)$ (average monthly spending, including zeros), the extensive margin $P(Y > 0)$ (participation in the category), and the intensive margin $E(Y | Y > 0)$ (spending conditional on participation). The unconditional mean is the primary metric, as it maps directly to the total-spending outcome in our main specifications. Since $E(Y) = P(Y > 0) \times E(Y | Y > 0)$, a category may exhibit some extensive-margin imbalance yet remain well balanced on the unconditional mean when the intensive margin offsets it.

The matched sample is well balanced. On the unconditional mean, the dimension corresponding to our DiD outcome, all 42 MCCs have $|\Delta_{\text{ND}}| < 0.05$, an order of magnitude below the standard threshold. A few categories, such as furniture, electronics, and drugs, show mild extensive-margin imbalance, with NDs in the 0.10–0.14 range that reflect small differences in participation rates. These do not translate into spending differences: for each such category the unconditional ND remains below 0.05, because the intensive margin offsets in the opposite direction. Balance therefore holds not only on aggregate covariates but across the full composition of consumer spending.

¹⁷The normalized difference is $\Delta_{\text{ND}} = (\bar{X}_T - \bar{X}_C) / \sqrt{S_T^2 + S_C^2}$, where $\bar{X}$ and S^2 denote group means and variances. Unlike the t -statistic, whose denominator shrinks mechanically with sample size, the ND denominator is invariant to N , so it measures the gap between groups relative to the natural cross-sectional variation in the variable rather than relative to sampling noise ([Imbens and Wooldridge, 2009](#); [Imbens and Rubin, 2015](#)).

Table B17: Pre-Period Normalized Differences across Merchant Categories

MCC	Category	Share	Unconditional $E(Y)$			Extensive $P(Y > 0)$			Intensive $E(Y Y > 0)$		
			Treat	Control	ND	Treat	Control	ND	Treat	Control	ND
Panel A: Risk and Behavioral Preferences											
6300	Insurance Company	1.19%	3.87	4.31	-0.009	3.7%	3.8%	-0.005	105.81	113.64	-0.04
5993	Tobacco	0.54%	2.84	2.90	-0.001	21.8%	18.7%	0.054	13.03	15.49	-0.021
6301	Insurance Agent	0.35%	0.68	0.77	-0.006	0.7%	0.8%	-0.006	92.39	95.68	-0.026
5813	Bars/Nightclubs	0.32%	1.10	1.14	-0.001	8.7%	9.5%	-0.020	12.68	11.99	0.008
7995	Lottery	0.29%	1.79	1.29	0.008	12.8%	9.6%	0.071	13.21	12.48	0.004
5921	Liquor	0.21%	0.52	0.49	0.001	5.5%	4.9%	0.019	9.35	10.00	-0.008
7990	Mahjong Parlor	<0.1%	0.04	0.06	-0.005	0.2%	0.1%	0.002	30.17	45.00	-0.112
5467	Health Supplements	<0.1%	0.14	0.13	0.002	1.3%	1.1%	0.013	10.92	12.31	-0.019
Panel B: Income-Sensitive Categories											
7011	Hotels	4.00%	19.45	21.17	-0.015	62.8%	58.9%	0.057	30.98	35.96	-0.034
6513	Property Management	2.19%	9.95	11.68	-0.019	3.1%	3.4%	-0.012	323.60	346.80	-0.030
7298	Beauty & SPA	1.50%	17.32	14.66	0.024	73.3%	64.6%	0.133*	23.64	22.68	0.007
8299	Education Services	1.33%	4.23	4.95	-0.012	4.8%	5.5%	-0.022	87.18	89.35	-0.009
5094	Jewelry	1.25%	4.98	4.82	0.003	25.8%	24.0%	0.030	19.30	20.13	-0.009
5691	Luxury Goods	1.09%	4.54	5.29	-0.019	26.4%	26.5%	-0.002	17.22	19.95	-0.038
8249	Vocational Training	0.93%	2.89	3.42	-0.011	2.3%	2.4%	-0.005	123.04	141.38	-0.048
4722	Travel Agency	0.34%	1.19	1.43	-0.007	2.3%	2.4%	-0.006	52.62	59.57	-0.021
4733	Attractions	0.29%	2.05	2.63	-0.028	11.2%	15.0%	-0.066	18.29	17.59	-0.020
5309	Duty-Free	0.27%	2.21	3.26	-0.026	3.5%	4.9%	-0.050	62.96	66.16	-0.018
4511	Airlines	0.07%	0.30	0.47	-0.010	3.1%	4.8%	-0.061	9.66	9.82	-0.002
Panel C: Inelastic Demand											
8062	Public Hospital	6.02%	23.38	25.53	-0.019	47.8%	49.9%	-0.030	48.91	51.11	-0.014
5122	Drugs	2.44%	15.27	14.10	0.013	82.6%	76.6%	0.106*	18.48	18.41	0.001

Continued on next page

Table B17 continued from previous page

MCC	Category	Share	Unconditional $E(Y)$			Extensive $P(Y > 0)$			Intensive $E(Y Y > 0)$		
			Treat	Control	ND	Treat	Control	ND	Treat	Control	ND
5912	Pharmacy	0.53%	1.93	2.08	-0.007	32.9%	32.8%	0.002	5.86	6.35	-0.013
9222	Fines	0.42%	1.53	1.90	-0.008	12.5%	13.5%	-0.021	12.24	14.08	-0.014
4900	Utilities	0.35%	1.26	1.37	-0.005	5.2%	5.3%	-0.001	24.04	25.89	-0.022
8011	Private Hospital	0.33%	0.90	0.91	0.000	4.2%	4.1%	0.007	21.14	22.34	-0.010
8021	Dental	0.18%	0.34	0.42	-0.005	1.3%	1.5%	-0.012	25.42	27.38	-0.005
9311	Tax Payment	0.16%	0.24	0.20	0.003	0.7%	0.7%	0.004	32.58	29.47	0.025
<i>Panel D: Digital Literacy Proxies</i>											
4121	Taxi	3.24%	23.74	28.42	-0.040	72.5%	69.6%	0.045	32.75	40.83	-0.060
4812	Mobile Devices	1.11%	4.27	3.87	0.009	15.0%	12.9%	0.043	28.46	30.03	-0.013
4115	Subway	0.28%	1.50	1.52	-0.002	29.1%	30.7%	-0.025	5.16	4.96	0.012
7523	Parking	0.23%	0.85	1.17	-0.015	18.6%	21.3%	-0.046	4.57	5.52	-0.020
4114	Highway Toll	0.13%	2.26	2.04	0.005	12.2%	11.1%	0.022	18.57	18.30	0.002
<i>Panel E: Non-Durable Consumption</i>											
5812	Restaurants	36.49%	248.51	265.41	-0.036	99.8%	99.5%	0.041	248.89	266.69	-0.038
5411	Supermarkets	17.41%	105.48	107.14	-0.007	99.5%	98.6%	0.066	105.99	108.63	-0.011
5699	Apparel & Accessories	4.42%	66.36	70.63	-0.016	93.1%	90.0%	0.079	71.26	78.48	-0.025
5422	Meat Provisioners	2.53%	33.43	33.88	-0.002	94.7%	89.4%	0.138*	35.30	37.88	-0.012
5423	Fruit Stores	0.15%	2.17	2.41	-0.012	43.7%	43.0%	0.010	4.96	5.60	-0.021
5942	Books	0.11%	0.58	0.81	-0.019	10.6%	14.4%	-0.081	5.44	5.65	-0.006
<i>Panel F: Durable Goods</i>											
5712	Furniture	5.46%	32.90	29.47	0.015	71.6%	62.4%	0.139*	45.94	47.21	-0.004
5732	Electronics	3.88%	25.97	25.20	0.004	71.4%	64.6%	0.103*	36.38	39.00	-0.011
5311	Department Stores	1.88%	15.47	20.28	-0.043	60.4%	59.7%	0.011	25.60	33.97	-0.059
5722	Appliances	0.86%	3.47	3.67	-0.004	12.4%	11.4%	0.022	27.99	32.18	-0.030

Continued on next page

Table B17 continued from previous page

MCC Category	Share	Unconditional $E(Y)$			Extensive $P(Y > 0)$			Intensive $E(Y Y > 0)$		
		Treat	Control	ND	Treat	Control	ND	Treat	Control	ND

Notes: Normalized differences $\Delta_{ND} = (\bar{X}_T - \bar{X}_C) / \sqrt{\bar{S}_T^2 + \bar{S}_C^2}$ for pre-period (2017–2019) monthly spending. Within each panel, MCCs are sorted by spending share (descending). Share is each MCC's percentage of total offline spending across the 42 categories (2018–2021). [Imbens and Wooldridge \(2009\)](#) note that linear regression is sensitive to specification when $|\Delta_{ND}| > 0.25$; $|\Delta_{ND}| < 0.10$ indicates excellent balance. Unconditional: mean yuan/month including zeros; Extensive: participation rate (%); Intensive: mean yuan/month conditional on $Y > 0$. * $|\Delta_{ND}| > 0.10$ on the extensive margin; the corresponding unconditional ND remains below 0.05 in all cases. Sample: 217,199 matched pairs (434,398 users). These 42 MCCs account for 53% of total offline spending in our sample; three catch-all categories (Miscellaneous Personal Services, Miscellaneous General Merchandise, and uncoded transactions) account for an additional 21%, and all remaining individual MCCs each represent less than 1.5%.

C Data and Variable Construction

This appendix describes the payment data used in our analysis, the classification of transactions into consumption categories, the representativeness of the sample relative to external benchmarks, and the construction of the Financial Sophistication Index used as auxiliary heterogeneity evidence.

Table C1: Spending Distribution and Credit Profile of the Matched Analysis Sample

	Mean	SD	P10	P25	P50	P75	P90	P99	% zero
<i>Panel A. Monthly spending and activity, 2019 (consumer-months)</i>									
Total spending (yuan)									
Treated	2,886	4,667	416	871	1,742	3,448	6,431	18,035	0.8
Control	2,879	17,263	402	865	1,746	3,455	6,405	17,640	1.0
In-person spending (yuan)									
Treated	1,955	3,110	109	401	1,013	2,240	4,556	14,898	1.9
Control	1,953	3,099	83	360	990	2,265	4,664	14,864	2.2
Online spending (yuan)									
Treated	931	3,349	9.9	118	382	999	2,250	8,258	9.5
Control	926	16,965	0	100	381	1,021	2,253	7,808	12.9
Number of transactions									
Treated	45.5	36.4	10	19	37	62	93	163	0.8
Control	42.9	35.8	8	17	34	59	89	159	1.0
<i>Panel B. Fintech credit, treated users, 2019</i>									
	Mean	Median	P90	N users					
Credit line (yuan)	9,938	7,467	22,392	216,762					
Revolving balance, all (yuan)	472	0	167	216,753					
Revolving balance, if positive (yuan)	6,401	3,400	16,700	24,774					
Share with positive revolving (%)	11.4	—	—	216,762					
Utilization, if positive (%)	53.5	55.0	90.4	24,774					

Note: Pre-shock (2019) statistics for the matched analysis sample of 217,199 pairs, complementing Table 1. Panel A pools consumer-months; “% zero” is the share of consumer-months with no spending (or no transaction). Panel B reports treated users’ fintech-credit profile averaged within 2019: the credit line, the monthly revolving balance (unconditional, and conditional on being positive), the share of users ever carrying a positive balance during 2019, and credit utilization (revolving balance divided by the credit line) conditional on a positive balance.

As we are using observed payments from Alipay as a proxy variable of consumers’ consumption expenditure, it is necessary to provide a detailed explanation of the composition of these payments, consumers’ usage patterns, and the criteria and procedures we applied in selecting the transaction data.

C.1 Payments for Consumption

There are three main categories of consumers' payment activities on Alipay: online payments, in-person payments, and transfers.

(1) *Online payments.*

So-called online payments include payments made via mobile devices or PC browsers, where merchants use Alipay's API to facilitate payments through the mobile app or web platform. After initiating a transaction, the consumer is redirected to the Alipay app or web platform to complete the payment process. Typical usage scenarios include e-commerce platforms, on-demand delivery services, ride-sharing services, and other digital transactions.

(2) *In-person payments*

In-person payments include scenarios where consumers are physically present during the transaction, using mobile devices or biometric authentication. When making a payment, consumers can scan the merchant's QR code or allow the merchant to scan their payment barcode if a scanner is available. With compatible point-of-sale equipments, consumers may also complete the payment through NFC contact or facial recognition. QR code payments are the most widely utilized among these options, primarily because merchants require only a plastic QR code card, which payment platforms typically provide for free.

(3) *Transfers*

Transfers involve consumers sending money to other Alipay accounts or bank account. This type of payment typically occurs between acquaintances serving purposes such as social gifting or splitting a bill. The main difference of transfer with the previous two categories is that consumers cannot use credit card or fintech credit in transfer.

For the first two payment types, recipients need proper authorization to receive payments. If the recipient is a business entity, they are assigned a Merchant Category Code (MCC) specific to their industry. Since digital payments are extensively utilized by small individual businesses, often informal and without formal registration, the platform develops algorithms to distinguish business transactions from personal ones.¹⁸ For these two types of transactions, we directly use the MCC labels or the algorithm's classification of long-tail merchants to classify transactions into their respective consumption categories.

For transfers, we decide based on whether the recipient is classified as a merchant by the algorithm and filtered out non-commercial transactions. Firstly, their share of the total payment volume is very small (less than 0.1%), and secondly, we have insufficient information on these


¹⁸The main goal behind creating the algorithms was to identify merchants and, based on their business activity, offer them micro-loans. We utilize the outcomes of this algorithm to filter consumption transactions. These algorithms analyze transaction network patterns, payment amounts, and transaction frequency. For merchants, the algorithms also predict the industry of the recipient.

payments.

Among the 411 MCC classifications, Table C2 lists the 17 MCC categories with transaction shares exceeding 1%. Restaurants top the list in terms of both transaction volume and count, followed by miscellaneous personal services and grocery stores, each accounting for the next largest share of transaction volume and frequency.

Figure C1 provides the monthly transaction count and the average amount per payment of consumers in the matched sample. Over the five years from 2017 to 2021, the data covers an average of 37.0 payments per month per consumer, with each payment averaging 76.0 *yuan* (\$11). As a comparison, the annual Diary of Consumer Payment Choice report conducted by the Federal Reserve recorded that a representative U.S. consumer made an average of 36.6 monthly payments during the same period¹⁹.

Figure C1: The Payment Usage of Consumers in the Matched Sample



Note: This figure shows the consumers' payment usage. Panel A plots the number of payments made by consumers in treatment and control groups from Jan. 2017 to Dec 2021. Panel B plots the average amount of each payment in the two groups.

C.2 Method of Payment

The methods of payment include credit card, debit card, Huabei, balance, and Yu'e Bao. Consumers can set up and rank their preferred methods or cards in advance. When a payment is initiated, the consumer's payment page displays a default payment method, which is the highest-ranked option that the merchant accepts. Alternatively, consumers can take an additional step to choose another accepted payment method.

¹⁹The survey is conducted in each October. In our sample, the average number of payments in each October is 39.6.




Figure C2: Usage across Different Payment Methods

Note: This figure shows the average share of amounts made with various payment methods in each group. The control group consumers use the three payment methods in roughly equal proportions. The treatment group consumers use fintech credit significantly higher than the other three methods, the share is over 60% in 2020.




Figure C3: Per-transaction Amount by Different Payment Methods

Note: This figure shows the average amounts per transaction made with various payment methods. The per-transaction amount for debit card payments is relatively large. In contrast, Fintech credit transactions are significantly lower than all other payment methods, indicating that it is widely used for small, everyday purchases that traditional payment systems do not cover.



Non-credit payment methods, including debit card, balance payment, and Yu'e Bao, are universally accepted for all payments. The use of credit-based payments, Huabei and credit cards, however, is subject to the merchant's acceptance. Some merchants may not accept Huabei or credit card payments because Alipay or the banks may impose withdrawal fees or transaction fees on these payment types. Some merchants may be unwilling to accept credit cards or fintech credit due to the additional processing fees, which generally range from 0.6% to 0.8% of the payment amount.

C.3 Construction of the Financial Sophistication Index

To better measure the financial literacy associated with fintech credit usage, we develop a Financial Sophistication Index utilizing two unique payment methods within Alipay. The two payment methods are Balance payment and Yu'e Bao, as briefly introduced in Section C.2.

Yu'e Bao is an innovative e-wallet linked to the money market fund. Consumers deposit money into it and can use it for payment at any time. At the same time, they earn a daily return from the money market fund, with annual interest rates ranging from 2% to 7%. In addition, Alipay also offers a traditional balance payment option, where consumers can transfer money into the balance and use it for payment, but it does not have interest. Figure C3 shows that the average per transaction amount of these two methods are similar, suggesting that the two

Figure C4: Distribution of Age and Digital Payment Adoption Time in Random Samples



Note: This figure shows the distribution of age and the time of using digital payment in the sample before matching. Figure A displays the distribution age, where we can observe that the Treatment (Huabei) group has a younger age distribution than the control group. Figure B shows the distribution of the time of using digital payment. The treatment group has a higher proportion of early users than the control group.

payment methods are equally convenient. However, since the returns from the money market fund can be considered risk-free, saving money put in Yu’e Bao is strictly dominant over saving money in a balance account and using it in the payment if consumers fully understand and accept the financial innovation. Ideally, rational consumers should opt for Yu’e Bao instead of using balance payments. However, we observed that only 27.4% of consumers did this in 2019. There are 27.2% consumers who strictly prefer using balance to Yu’e Bao.

For each consumer i , we define the financial sophistication index as the share of “interest generated balance” used in all transactions in 2019:

$$FS_i = MMF_i / (MMF_i + Balance_i) \quad (C1)$$




Figure C5: Distribution of Financial Sophistication Index

Note: This figure shows the distribution of the Financial Sophistication Index in both the control and treatment groups. The distribution is concentrated at both ends (around 0 and 1). Overall, the FSI for the treatment group is higher than that of the control group.

Therefore, $FS = 1$ means very sophisticated, and the consumer is actively taking advantage of the innovative financial payment method. And $FS = 0$ means unsophisticated. Our interpretation in this situation is that these consumers either lack awareness of the interest advantages provided by the payment method, or they are aware but choose not to use this innovative payment method due to distrust or other conservative considerations.

Among the 77.8% consumers who have a well-defined FS index in 2019, Figure C5 plots the distribution of the index of two groups. The treatment group consumers are more sophisticated than control group under this measure. However, in both groups we found that the index shows

a higher proportion at the extremes of 0 and 1, while the middle range displays a relatively even distribution.




Figure C6: Distribution of Financial Sophistication Index by Education

Note: This plot shows the distribution of the Financial Sophistication Index in each education group. First, the high education group has a larger share of consumers with $FS = 1$. However, in each education group, the distribution from 0 to 1 is similar, showing that the FS index and education are not completely correlated.

As shown in Figure C6, the high-education group has a larger share of financially sophisticated consumers. In consumers with a doctoral or master's degree, the proportion of $FS = 1$ is 46% and 40%, respectively, while 30% in the high school or below group. On the other hand, even in high education groups, a significant amount of consumers are not financially sophisticated, suggesting the Financial Sophistication Index and educational level are not entirely aligned. Therefore, our newly defined index highlights the consumer's understanding and acceptance of different payment methods within the fintech payment platform, offering a new dimension in measuring financial literacy.

Table C2: Top MCC Merchants Types in Payments

No.	MCC code	Amount	Count	Group	Category	Merchant Type
1	5812	16.8%	25.3%	SND	Food Away From Home	Eating places and Restaurants
2	7299	8.8%	4.9%	ND	Misc. Nondurables	Miscellaneous Personal Services
3	5411	7.9%	11.9%	SND	Groceries	Grocery Stores, Supermarkets
4	5699	4.4%	1.5%	ND	Retail Nondurables	Misc. Apparel and Accessory Shops
5	5399	3.2%	2.2%	ND	Retail Nondurables	Misc. General Merchandise
6	8062	2.6%	0.8%	ND	Medical Copay	Hospitals
7	5422	2.5%	5.4%	SND	Groceries	Meat Provisioners - Freezer and Locker
8	5712	2.5%	0.6%	D	Home Improvement	Furniture, Home Furnishings, and Equipment Stores
9	5311	1.9%	0.9%	ND	Department Stores	Department Stores
10	5732	1.8%	0.5%	D	Retail Durables	Electronic Sales
11	7011	1.8%	0.4%	D	Hotels & Rental Cars	Lodging - Hotels, Motels, Resorts, Central Reservation Services
12	4121	1.6%	5.0%	SND	Ground Transportation	Taxis and Limousines
13	7994	1.6%	1.1%	D	Entertainment	Video Game Arcades/Establishments
14	7531	1.4%	0.5%	D	Auto Repair	Automotive Body Repair Shops
15	7298	1.4%	0.5%	SND	Professional & Personal Services	Health and Beauty Shops
16	5331	1.2%	2.7%	ND	Retail Nondurables	Variety Stores
17	5122	1.1%	0.7%	ND	Drug Stores	Drugs, Drug Proprietors, and Druggist's Sundries

Notes: This table shows merchant types with a transaction amount share greater than 1% of all payments. Among all payments records, 8.879% of the transaction amount (5.814% of the transaction count) was directed to recipients without an assigned MCC code. Out of 407 MCC codes, we successfully matched 369, and 0.018% of the total transaction amount was unmatched (0.016% of the transaction count).

	Age	Big City	Education	Consumption	FS Index	Treatment
Age	1	-0.025	-0.466	0.198	-0.032	-0.018
Big city		1	0.127	0.103	0.02	0.008
Education			1	-0.109	0.062	-0.046
Consumption				1	0.024	0.019
FS Index					1	0.164
Treatment						1

Table C3: Correlation of Variables Used in Heterogeneity Analysis

Note: This table shows the correlation between the variables used in heterogeneity analysis. Education is denoted as 1 to 4, to represent the four levels of education: high school education or below, post-secondary, college, and those with graduate degrees. FS index refers to the Financial Sophistication Index.

D External Validity and Sensitivity Analyses

This appendix establishes the external validity of our baseline estimates. It demonstrates that they are robust to (and in fact conservative relative to) a range of alternative estimators applied to the full 3.5-million-consumer population.

D.1 IPW, PSM, and reweighting estimators

Our baseline analysis matches 500,000 Huabei (BNPL) users to non-users using 1:1 nearest-neighbor matching with Mahalanobis distance within propensity score calipers, with exact matching on gender, region, and platform entry time. The matching succeeds for 217,199 treated users (43.4%). The remaining $\sim 283,000$ treated observations are dropped from the matched sample. This raises a natural concern: *are the estimated treatment effects representative of the broader Huabei user population?*

We address this concern through three complementary strategies that together demonstrate the robustness and external validity of our findings:

- (i) **Inverse Probability Weighting (IPW) estimation** on the full sample of 3.5 million consumers, imposing no matching restriction;
- (ii) **Many-to-one propensity score matching** with relaxed calipers to achieve higher matching rates;
- (iii) **Reweighting the matched sample** to restore representativeness of the full treatment population.

All analyses use the same outcome variables and time windows as the baseline specification. The estimation sample comprises 3,497,126 consumers observed over 60 months (January 2017–December 2021), yielding a panel of approximately 210 million person-month observations.

D.1.1 Data and Variable Definitions

The full analysis sample consists of 499,835 randomly sampled Huabei users (treatment) who activated Huabei before January 2020 and 2,997,291 randomly sampled Alipay non-users (control), for a total of 3,497,126 consumers observed over 60 months (209,827,560 person-month observations).

We use three outcome variables: $Y_{it}^{\text{level}} = \text{TotalPay}_{it}$ (consumption in levels for PPML estimation, our primary specification); $Y_{it}^{\text{LPM}} = \mathbf{1}[\text{TotalPay}_{it} > 0]$ (extensive margin); and $Y_{it}^{\text{int}} = \log(\text{TotalPay}_{it}) \mid \text{TotalPay}_{it} > 0$ (intensive margin).

We define five post-treatment windows, all sharing the same 24-month pre-period (months 13–36, corresponding to January 2018–December 2019):

Label	Post-period months	Calendar dates
Three months	38–40	February–April 2020
Six months (H1)	37–42	January–June 2020
Year 2020	37–48	January–December 2020
Year 2021	49–60	January–December 2021
Two years	37–60	January 2020–December 2021

The shock date is January 2020 (month 37), when the COVID-19 pandemic created a negative income shock that differentially affected consumption smoothing across credit-access groups.

In all strategies, the propensity score is estimated using the same set of pre-treatment covariates: log quarterly consumption ($\ln C_{i,2019Q1}$, $\ln C_{i,2019Q2}$, $\ln C_{i,2019Q3}$, $\ln C_{i,2019Q4}$, $\ln C_{i,Dec2018}$), age and age squared, gender (factor), platform entry time (factor, quarter of Alipay account creation), and region (factor, 70 groups). Missing consumption values are imputed as zero before taking logs.

D.1.2 Specification 1: IPW Estimation on the Full Sample

We estimate the propensity of Huabei adoption using a logistic regression on the full sample of 3,497,126 consumers:

$$\text{logit}[\hat{e}(X_i)] = \alpha + \sum_{k=1}^5 \beta_k \ln C_{i,k} + \gamma_1 \text{Age}_i + \gamma_2 \text{Age}_i^2 + \text{FE}_{\text{gender}} + \text{FE}_{\text{enter_time}} + \text{FE}_{\text{region}} \quad (\text{D1})$$

where $\hat{e}(X_i) = \Pr(\text{Huabei}_i = 1 \mid X_i)$ is the estimated propensity score. This specification mirrors the exact-match and distance variables used in the baseline matching procedure. Following [Crump et al. \(2009\)](#), we trim observations with extreme propensity scores below the 2nd percentile or above the 98th percentile of the treated group’s distribution; after trimming, 2,987,387 of 3,497,126 observations (85.4%) are retained.

Method 1C: IPW-Weighted Panel Fixed Effects To ensure comparability with the baseline panel specification, we estimate the same two-way fixed effects model with IPW weights on the full 3.5M trimmed panel (179,243,220 person-month observations):

$$Y_{it} = \beta \cdot (\text{Huabei}_i \times \text{Post}_t) + \alpha_i + \delta_t + \varepsilon_{it} \quad (\text{D2})$$

where α_i are individual fixed effects, δ_t are month fixed effects, and each control observation is weighted by $w_i = \hat{e}(X_i)/(1 - \hat{e}(X_i))$ (treated units receive weight 1); control weights are normalized so that their mean equals one. For each window, we estimate two specifications: an LPM with $Y_{it} = \mathbf{1}[\text{TotalPay}_{it} > 0]$ for the extensive margin, and $Y_{it} = \log(\text{TotalPay}_{it})$ on the subsample $\text{TotalPay}_{it} > 0$ for the intensive margin. Standard errors are heteroskedasticity-robust, equivalent to individual-clustered SEs given that the regressor varies at the individual $\times$ period level.

Method 1D: Collapsed PPML Panel PPML with 3 million individual fixed effects is impractical at this scale, so we estimate a collapsed cross-sectional PPML (Silva and Tenreiro, 2006). For each time window, we collapse to person-level post-period means $\bar{Y}_{i,\text{post}} = \frac{1}{T_{\text{post}}} \sum_{t \in \text{post}} \text{TotalPay}_{it}$ and estimate:

$$E[\bar{Y}_{i,\text{post}} | X_i] = \exp(\alpha + \beta \cdot \text{Huabei}_i + \gamma \cdot \bar{Y}_{i,\text{pre}}) \quad (\text{D3})$$

by quasi-Poisson GLM with IPW weights. We control for mean pre-period consumption as the collapsed analogue of individual fixed effects; β is interpreted as $\exp(\beta) - 1 =$ percentage change in consumption. Standard errors are from the sandwich estimator with the IPW weights.

Table D1: Specification 1: IPW Estimates on Full 3.5M Sample (Pre = months 24–36)

Window	PPML (%)		LPM (pp)	
	Coef	Region SE	Coef	Region SE
Three months	14.91***	[0.011]	6.16***	[0.26]
Six months	13.03***	[0.010]	5.19***	[0.22]
Year 2020	14.34***	[0.012]	4.64***	[0.22]
Two years	14.75***	[0.014]	4.93***	[0.25]

Notes: PPML estimated by collapsed quasi-Poisson GLM controlling for pre-period mean consumption; the LPM regressions are panel FE on the IPW-trimmed 3.5M panel with individual and month fixed effects. All regressions weighted by the IPW weights described in Method 1C. PPML coefficients reported as $(\exp(\hat{\beta}) - 1) \times 100$; LPM coefficients in percentage points. Region-clustered standard errors (66 regions) in square brackets. *, **, *** denote significance at the 10%, 5%, and 1% levels. $N_{\text{individuals}} = 2,987,387$.

D.1.3 Specification 2: Many-to-One Propensity Score Matching

We show that the treatment effect is robust to alternative matching specifications that potentially achieve higher matching rates. While our baseline uses Mahalanobis distance with exact matching constraints, here we employ K-nearest-neighbor propensity score matching with varying ratios and caliper widths. For each specification, we match each treated unit to its K nearest

control units within a caliper of $c \times \text{SD}(\hat{e})$, with matching done *with replacement*. We consider three specifications: Spec A (1:1, caliper $0.10 \times \text{SD}$), Spec B (1:3, $0.25 \times \text{SD}$), and Spec C (1:5, $0.50 \times \text{SD}$). Each selected control j receives weight $1/K_i^*$ per match (with $K_i^* \leq K$ being the number of matched controls), and these weights are summed and normalized across treated units so that the mean control weight equals 1.

For each matched sample we estimate three regressions: PPML, LPM for the extensive margin, and $\log|Y| > 0$ for the intensive margin on the positive-outcome subsample, all with individual and month fixed effects and matching weights. We also report an LPM specification with city×month fixed effects as a robustness check. Standard errors are heteroskedasticity-robust.

Table D2: Specification 2: PSM Matching Robustness Across Post-Windows (Pre = months 24–36, Panel FE)

Spec	N indiv	PPML semi-elasticity ($e^{\hat{\beta}} - 1$, %)			
		3 months	6 months	Year 2020	Two years
A: 1:1 ($c = 0.10 \text{ SD}$)	902,618	19.98*** [0.009]	18.14*** [0.008]	19.99*** [0.010]	20.93*** [0.011]
B: 1:3 ($c = 0.25 \text{ SD}$)	1,459,240	19.75*** [0.008]	18.20*** [0.008]	20.13*** [0.010]	21.06*** [0.011]
C: 1:5 ($c = 0.50 \text{ SD}$)	1,826,294	19.79*** [0.008]	18.18*** [0.008]	20.09*** [0.010]	20.99*** [0.011]
<i>Reference: Baseline (434K matched, panel FE + city×month):</i>					
Baseline		9.03*** [0.009]	7.76*** [0.008]	8.46*** [0.009]	8.56*** [0.010]

Notes: PSM specifications use panel `fepois` with individual and month fixed effects on each matched sample with matching weights $1/K$ applied to controls; baseline uses individual + city×month fixed effects on the 434K Mahalanobis-matched sample. Pre-period months 24–36. Region-clustered SEs (66 regions) in square brackets. *, **, *** denote significance at 10%, 5%, 1%. Matching is K-nearest-neighbor PSM with replacement on the full 3.5M sample; 100% match rate for all three calipers indicates all 499,835 treated users found at least one control within caliper.

D.1.4 Specification 3: Reweighting the Matched Sample

For analyses that require variables available only in the matched sample (e.g., heterogeneity by education, occupation, external credit), we reweight matched treatment observations to restore representativeness of the full 500K treatment population, following [Stuart et al. \(2011\)](#) and [Hartman et al. \(2015\)](#). Within the 500K treatment group, we first estimate the probability of being

successfully matched:

$$\text{logit}[\hat{\pi}_i] = \alpha + \sum_{k=1}^5 \beta_k \ln C_{i,k} + \gamma_1 \text{Age}_i + \gamma_2 \text{Age}_i^2 + \delta_1 \text{CV}_i + \delta_2 \text{Growth}_i + \text{FE}_{\text{gender}} + \text{FE}_{\text{enter_time}} + \text{FE}_{\text{region}} \quad (\text{D4})$$

where CV_i is the coefficient of variation of quarterly consumption in 2019 and Growth_i is year-on-year consumption growth (2019 vs. 2018). The external validity weight is $w_i^{\text{ext}} = 1/\hat{\pi}_i$, capped at the 99th percentile and normalized to have mean 1 across the matched treatment sample.

The 217,199 matched treated users differ from the full 500K treatment population mainly in three dimensions: matched consumers have approximately 8% higher annual consumption (moderate, stable spenders find comparable controls more easily); lower consumption volatility (volatile spenders lack close control counterparts); and a different gender composition reflecting the exact-match constraint on gender.

We run the same panel FE specifications as the baseline on two samples unweighted (the baseline 217,199 matched pairs $\times$ 60 months, subset to relevant windows) and reweighted (same sample with weights w_i^{ext}) using three regression types: PPML, LPM for the extensive margin, and intensive margin $\log(Y)|Y > 0$ on the positive-outcome subsample.

Table D3: Specification 3: Reweighted vs. Unweighted Matched Samples (434K, Pre = months 24–36)

Window	434K Unweighted			434K Reweighted		
	PPML%	LPM (pp)	$\log Y > 0$	PPML%	LPM (pp)	$\log Y > 0$
3 months	9.03*** [0.009]	1.76*** [0.14]	0.123*** [0.008]	10.05*** [0.011]	2.95*** [0.26]	0.171*** [0.013]
6 months	7.76*** [0.008]	1.32*** [0.11]	0.121*** [0.007]	8.24*** [0.009]	2.59*** [0.26]	0.163*** [0.011]
Year 2020	8.46*** [0.009]	1.20*** [0.12]	0.132*** [0.009]	8.95*** [0.010]	2.46*** [0.26]	0.176*** [0.012]
Two years	8.56*** [0.010]	1.41*** [0.17]	0.127*** [0.009]	8.60*** [0.011]	2.45*** [0.28]	0.166*** [0.012]

Notes: All specifications use panel `fepois`/`feols` with individual and city $\times$ month fixed effects on the 434K matched sample; pre-period months 24–36. LPM in percentage points; PPML as $(\exp(\hat{\beta}) - 1) \times 100$; intensive margin estimated as $\log(Y) | Y > 0$ on the positive-outcome subsample. Region-clustered SEs (66 regions) in square brackets. *, **, *** denote significance at 10%, 5%, 1%. 434K sample: 217,199 matched treated users + matched controls (16.1M person-month obs for the two-year window). Reweighted estimates use external validity weights from eq. (D4).

Table D4: Two-Year ATT Across All Methods and Samples (Pre = months 24–36)

Method	PPML (%)	LPM (pp)
<i>Baseline:</i>		
Panel FE, 434K matched	8.56*** [0.010]	1.41*** [0.17]
<i>Specification 1: IPW on full sample (3.5M):</i>		
IPW (collapsed PPML)	14.75*** [0.014]	4.93*** [0.25]
<i>Specification 2: PSM matching (panel FE):</i>		
1:1, caliper = 0.10 SD	20.93*** [0.011]	4.96*** [0.26]
1:3, caliper = 0.25 SD	21.06*** [0.011]	4.96*** [0.25]
1:5, caliper = 0.50 SD	20.99*** [0.011]	4.96*** [0.26]
<i>Specification 3: Reweighted matched (panel FE):</i>		
434K unweighted	8.56*** [0.010]	1.41*** [0.17]
434K reweighted	8.60*** [0.011]	2.45*** [0.28]

Notes: All estimates use the two-year post-period (2020–2021) with a pre-period of months 24–36 (December 2018 + 2019). Region-clustered SEs (66 regions) in square brackets. *, **, *** denote significance at 10%, 5%, 1%. Base-line and Specification 3 use panel `fepois/feols` with individual + city×month fixed effects on the 434K matched panel; Specification 2 uses panel `fepois` with individual + month fixed effects on each PSM-matched panel ($N_A = 902,618$, $N_B = 1,459,240$, $N_C = 1,826,294$); Specification 1 uses collapsed quasi-Poisson with pre-period mean consumption control on the IPW-weighted 3.5M sample ($N = 2,987,387$). PPML coefficients reported as $(\exp(\hat{\beta}) - 1) \times 100$; LPM in percentage points.

D.1.5 Comprehensive Comparison Across All Methods

D.1.6 Discussion

The full-sample PPML estimates (14.8% for IPW, $\approx 21\%$ for PSM) are larger than the baseline matched-sample PPML estimate (8.6%). This gap is consistent with the composition of the unmatched group: treated consumers dropped from the baseline analysis tend to have more extreme or volatile consumption patterns (high CV, extreme growth rates), which makes them harder to match and, plausibly, more responsive to credit access during income shocks. Under this reading, the baseline estimate is conservative relative to the broader treated population, although we cannot rule out that part of the gap reflects weaker overlap and heavier reliance on the propensity-score model in the full-sample estimators.

The consistency across different estimation strategies supports robustness. PSM with varying K (1, 3, 5) and calipers (0.10, 0.25, 0.50 SD) produces nearly identical PPML estimates (20.9%–21.1% at two years) and LPM estimates (≈ 5.0 pp), all highly significant with region-clustered SEs. Reweighting the 434K matched sample toward the full treatment population leaves the PPML estimate essentially unchanged at the two-year horizon (8.6% vs. 8.6%) while raising the LPM from 1.41 pp to 2.45 pp, indicating that the matching procedure does not generate selection on the observable PPML margin but does compress the extensive-margin response.

The PPML estimates show some variation across methods (8.6%–21.1%). The IPW PPML (14.8%) uses collapsed quasi-Poisson on the full 3.5M sample (panel PPML with 3 million individual fixed effects is computationally impractical, and the collapsed estimator controls for pre-period mean consumption as the cross-sectional analogue of individual fixed effects), while PSM PPML (20.9%–21.1%) uses panel Poisson with individual fixed effects on matched samples; the difference reflects both the estimation approach (collapsed vs. panel) and the sample-composition effect described above. The LPM estimates show tighter convergence across methods.

D.1.7 Dynamic Difference-in-Differences

To examine the time path of treatment effects and assess parallel pre-trends, we estimate a dynamic (event study) specification on the 434K baseline matched sample:

$$Y_{it} = \sum_{\substack{k=-24 \\ k \neq -1}}^{23} \beta_k \cdot (\text{Huabei}_i \times \mathbf{1}[t = k]) + \alpha_i + \delta_t + \varepsilon_{it} \quad (\text{D5})$$

where event time $k = 0$ corresponds to January 2020 (the treatment date) and $k = -1$ (December 2019) is the reference period. We estimate (D5) for three outcome variables: PPML on consumption levels, LPM for the extensive margin, and $\log|Y| > 0$ for the intensive margin on the

positive-outcome subsample.

We report three levels of clustering: individual-clustered SEs (the natural choice since treatment is assigned at the individual level); city-clustered SEs (328 cities, accounting for spatial correlation within cities); and region-clustered SEs (66 regions, the most conservative). City-clustered SEs are approximately 2.5–3× larger than individual-clustered SEs, and region-clustered SEs are 5–6× larger. All coefficients remain statistically significant at the 0.1% level under all three clustering levels; with 66 region clusters, the two-year static DiD t-statistics are 8.6 (PPML) and 11.1 (LPM).

D.2 Sensitivity to pre-trend violations

Table D5 reports Rambachan and Roth (2023) sensitivity bounds for the quarterly event-study estimates of the treatment effect on total consumption. The relative-magnitude exercise (Panel A) takes the largest pre-period deviation as a benchmark and asks whether post-period estimates remain informative when post-period violations of parallel trends are allowed to be as large as a multiple $\bar{M}$ of that benchmark. The acute-quarter effect remains positive for violations up to 1.5 times the benchmark, and the year-one average up to half the benchmark; longer-horizon averages are more sensitive, as confidence sets widen mechanically with horizon. Panel B reports the corresponding smoothness-restriction bounds, which allow the differential trend to change slope by at most M per quarter.

Two features make this exercise conservative in our setting. First, the benchmark deviations arise in 2018, before the window on which the sample is matched; they reflect pre-match level differences that the matching subsequently absorbs rather than a trend at the shock date. Second, within the identification window (December 2018 through December 2019) the quarterly pre-period coefficients are jointly indistinguishable from zero (Wald $p = 0.42$) and the estimated linear differential trend is 0.10% per month ($p = 0.35$), so the local-to-the-shock evidence gives no indication of an ongoing violation.

Table D5: Sensitivity of Event-Study Estimates to Pre-Trend Violations

	Original	Robust confidence set, relative magnitudes $\tilde{M}$			
	95% CI	$\tilde{M} = 0.5$	$\tilde{M} = 1$	$\tilde{M} = 1.5$	$\tilde{M} = 2$
<i>Panel A. Rambachan–Roth relative-magnitude bounds (Δ^{RM})</i>					
2020Q1 (acute quarter)	[7.0, 10.0]	[5.6, 11.1]	[3.9, 12.8]	[2.0, 14.5]	[-0.1, 15.6]
Year-1 average (2020)	[6.4, 9.5]	[2.3, 13.2]	[-2.6, 15.8]	[-7.5, 15.8]	[-12.4, 15.8]
Year-2 average (2021)	[6.1, 10.4]	[-5.3, 21.4]	[-18.1, 21.6]	[-21.6, 21.6]	[-21.6, 21.6]
Two-year average	[6.3, 9.9]	[-1.6, 17.3]	[-10.3, 18.2]	[-18.2, 18.2]	[-18.2, 18.2]
<i>Panel B. Smoothness-restriction bounds (Δ^{SD}), slope change M per quarter</i>					
		$M = 0.01$	$M = 0.03$	$M = 0.05$	
2020Q1 (acute quarter)	[7.0, 10.0]	[6.0, 11.0]	[4.2, 12.9]	[2.2, 14.9]	
Year-1 average (2020)	[6.4, 9.5]	[1.5, 14.6]	[-8.5, 24.6]	[-18.5, 34.6]	
Year-2 average (2021)	[6.1, 10.4]	[-19.4, 36.4]	[-69.3, 86.3]	[-119.3, 136.3]	
Two-year average	[6.3, 9.9]	[-8.9, 25.5]	[-38.9, 55.5]	[-68.7, 85.3]	

Note: Rambachan and Roth (2023) robust confidence sets for the quarterly PPML event-study coefficients on total consumption (coefficients approximately equal percentage effects; entries multiplied by 100). Panel A allows post-period violations of parallel trends up to $\tilde{M}$ times the largest pre-period deviation (relative-magnitudes class Δ^{RM}); Panel B allows the differential trend to change slope by at most M per quarter (smoothness class Δ^{SD}). The original 95% confidence intervals cluster standard errors by region. The pre-period deviations that define the benchmarks arise in 2018, before the matched window (quarterly deviations relative to 2019Q4: -3.3%, -4.4%, -4.5%, -1.9% for 2018Q1–Q4); within the identification window (December 2018 through December 2019) the pre-coefficients are jointly indistinguishable from zero (Wald $p = 0.42$).

E Decomposition of the Consumption Response

This appendix provides the full set of decompositions underlying the channel, frequency, and margin results reported in Section 2. We report PPML and LPM estimates for four post-treatment windows (3 months, 6 months, 1 year, 2 years), separately for total, offline, and online consumption, and both for the 434K unweighted matched sample and for the externally reweighted sample constructed in Appendix D.1. Dynamic event-study counterparts are reported for each decomposition.

This appendix reports the full set of decompositions behind the margin, channel, and frequency results summarized in Table 4 (Panels B–D), together with their dynamic counterparts. The sample, estimation windows, specifications (PPML for amounts and counts, LPM for the extensive margin, $\log(Y) \mid Y > 0$ for the intensive margin), fixed effects, and region-clustered standard errors are identical to those defined in Sections 1.2 and 2 and are not restated here. Each table reports the unweighted 434K matched sample (“Unwtd”) and the same sample reweighted by the external-validity weights of Appendix D.1 (“Rewtd”). Following Chen and Roth (2024), the extensive/intensive split is a descriptive decomposition: the conditional-on-positive sample is itself affected by treatment, so causal inference rests on the PPML estimates in levels.

E.1 Spending amounts by channel

Table E1: Treatment Effects on Spending Amounts (PPML)

Window	Total Amount		Offline Amount		Online Amount	
	Unwtd	Rewtd	Unwtd	Rewtd	Unwtd	Rewtd
3 months	0.086*** [0.009]	0.096*** [0.011]	0.140*** [0.010]	0.155*** [0.013]	−0.021* [0.012]	−0.026** [0.012]
6 months	0.075*** [0.008]	0.079*** [0.009]	0.116*** [0.010]	0.122*** [0.011]	−0.009 [0.012]	−0.011 [0.013]
1 year	0.081*** [0.009]	0.086*** [0.010]	0.120*** [0.012]	0.124*** [0.013]	−0.001 [0.012]	−0.002 [0.012]
2 years	0.082*** [0.010]	0.082*** [0.011]	0.122*** [0.013]	0.121*** [0.014]	−0.003 [0.012]	−0.006 [0.012]
<i>N</i> (2yr)	16,065,437		16,065,391		16,006,311	

Notes: PPML (fepois) estimates with individual and city×month fixed effects; pre-period months 24–36 (December 2018 + 2019). The coefficient $\hat{\beta}$ is a semi-elasticity; $e^{\hat{\beta}} - 1$ gives the percentage effect. Region-clustered standard errors (66 regions) in square brackets. “Unwtd” = unweighted 434K matched sample; “Rewtd” = reweighted with external validity weights. *, **, *** denote significance at the 10%, 5%, and 1% levels with region-clustered SEs. All total and offline amount coefficients are significant at $p < 0.001$; online amount coefficients are small and statistically insignificant outside the acute 3-month window.

Huabei access raises total spending by 8–10%, driven entirely by in-person spending (12–17%); online spending is unchanged at every horizon beyond the acute window. This is the appendix-level counterpart of Table 4, Panel C, and the basis for the in-person concentration discussed in Section 4.3.

Table E2: Treatment Effects on Transaction Counts (PPML)

Window	Total Count		Offline Count		Online Count	
	Unwtd	Rewtd	Unwtd	Rewtd	Unwtd	Rewtd
3 months	0.015** [0.006]	0.031*** [0.007]	0.078*** [0.007]	0.114*** [0.009]	-0.050*** [0.004]	-0.041*** [0.005]
6 months	0.028*** [0.005]	0.046*** [0.005]	0.082*** [0.007]	0.116*** [0.008]	-0.036*** [0.004]	-0.024*** [0.004]
1 year	0.029*** [0.005]	0.055*** [0.007]	0.068*** [0.007]	0.108*** [0.009]	-0.025*** [0.004]	-0.008* [0.005]
2 years	0.016*** [0.006]	0.044*** [0.007]	0.052*** [0.008]	0.093*** [0.009]	-0.034*** [0.003]	-0.016*** [0.005]
<i>N</i> (2yr)	16,065,437		16,065,391		16,006,311	

Notes: PPML (fepois) estimates with individual and city×month fixed effects; pre-period months 24–36. Region-clustered SEs (66 regions) in square brackets. *, **, *** denote significance at the 10%, 5%, and 1% levels. Offline count effects are positive and significant at $p < 0.001$ for every window/sample. Online count effects are *negative* and (with one exception in the reweighted 1-year cell) statistically significant, indicating fewer online orders among treated consumers.

E.2 Transaction frequency

In-person transaction counts rise by 5–12% while online order counts fall by 1–5%; combined with unchanged online amounts, the remaining online orders are larger on average. Treated consumers purchase in person more often, not only in larger amounts.

E.3 Extensive and intensive margins

The probability of any purchase rises by 1.2–1.8 pp (unweighted), three to seven times larger in person than online. Conditional on purchasing, total spending rises by 0.12–0.18 log points and in-person spending by 0.18–0.26 log points; the online intensive margin is slightly negative. As noted above, the conditional-on-positive estimates are descriptive.

E.4 Summary and reweighted sample

Across all margins the response is concentrated in person: amounts (+12.9%), frequency (+5.3%), and the extensive margin (+3.4 pp) all rise, while online amounts are flat and online order counts decline. The reweighted estimates are uniformly larger by roughly 30–60%, because the external-validity weights up-weight hard-to-match consumers with stronger responses; sign patterns and significance are unchanged, so the decomposition is not an artifact of the matching procedure.

Table E3: Treatment Effects on Extensive Margin: Any Purchase (LPM)

Window	Any Purchase		Any Online		Any Offline	
	Unwtd	Rewtd	Unwtd	Rewtd	Unwtd	Rewtd
3 months	1.76 pp*** [0.14]	2.95 pp*** [0.26]	0.85 pp*** [0.11]	1.90 pp*** [0.18]	4.57 pp*** [0.28]	6.30 pp*** [0.46]
6 months	1.32 pp*** [0.11]	2.59 pp*** [0.26]	0.67 pp*** [0.09]	1.88 pp*** [0.19]	3.72 pp*** [0.25]	5.42 pp*** [0.43]
1 year	1.20 pp*** [0.12]	2.46 pp*** [0.26]	0.57 pp*** [0.10]	1.95 pp*** [0.21]	3.17 pp*** [0.25]	4.90 pp*** [0.43]
2 years	1.41 pp*** [0.17]	2.45 pp*** [0.28]	0.44 pp*** [0.12]	1.77 pp*** [0.22]	3.37 pp*** [0.31]	4.89 pp*** [0.46]
<i>N</i> (2yr)	16,065,437					

Notes: Linear probability model with individual and city×month fixed effects; pre-period months 24–36. Coefficients and standard errors reported in percentage points (pp). Region-clustered SEs (66 regions) in square brackets. *, **, *** denote significance at the 10%, 5%, and 1% levels. All coefficients are significant at $p < 0.001$ (minimum $t = 3.6$ for the online column, 2-year window).

Table E4: Treatment Effects on Intensive Margin: $\log(Y) \mid Y > 0$

Window	Total		Offline		Online	
	Unwtd	Rewtd	Unwtd	Rewtd	Unwtd	Rewtd
3 months	0.123*** [0.008]	0.171*** [0.013]	0.177*** [0.010]	0.248*** [0.014]	−0.026*** [0.004]	−0.005 [0.006]
6 months	0.121*** [0.007]	0.163*** [0.011]	0.181*** [0.009]	0.245*** [0.013]	−0.017*** [0.004]	−0.002 [0.005]
1 year	0.132*** [0.009]	0.176*** [0.012]	0.196*** [0.012]	0.259*** [0.015]	−0.010*** [0.003]	0.003 [0.004]
2 years	0.127*** [0.009]	0.166*** [0.012]	0.198*** [0.012]	0.256*** [0.015]	−0.026*** [0.004]	−0.014*** [0.005]
<i>N</i> (2yr)	15,615,783		14,802,509		14,082,238	

Notes: Panel fixed-effects (feols) estimates of $\log(Y) \mid Y > 0$ on the positive-outcome subsample, with individual and city×month fixed effects; pre-period months 24–36 on the 434K matched sample. Region-clustered SEs (66 regions) in square brackets. *, **, *** denote significance at the 10%, 5%, and 1% levels. Total and offline log-intensive effects are significant at $p < 0.001$ in every cell. Online effects in the unweighted sample are *negative* and significant; the reweighted online cells are essentially zero except at the 2-year horizon.

Table E5: Summary: 2-Year Treatment Effects Across Channels (Unweighted)

Outcome	Model	Total	Offline	Online
<i>Panel A: Spending Amounts</i>				
Amount (PPML)	fepois	8.6%*** [0.010]	12.9%*** [0.013]	−0.3% [0.012]
<i>Panel B: Transaction Frequency</i>				
Count (PPML)	fepois	1.6%*** [0.006]	5.3%*** [0.008]	−3.4%*** [0.003]
<i>Panel C: Extensive Margin</i>				
Any purchase (LPM)	feols	1.41 pp*** [0.17]	3.37 pp*** [0.31]	0.44 pp*** [0.12]
<i>Panel D: Intensive Margin</i>				
$\log(Y) \mid Y > 0$	feols	0.127*** [0.009]	0.198*** [0.012]	−0.026*** [0.004]

Notes: All estimates use individual + city×month FE on the 434K unweighted matched sample, 2-year window, pre-period months 24–36. Region-clustered SEs (66 regions) in brackets. PPML effects converted to percent ($e^{\hat{\beta}} - 1$). LPM in percentage points. *** $p < 0.001$ with region-clustered SEs.

E.5 Event-study estimates

Figures E1 and E2 report the dynamic counterparts for transaction counts (PPML) and the extensive margin (LPM) by channel, estimated on the canonical window December 2018 through December 2021 with December 2019 as the reference month. Pre-period coefficients are small with no systematic trend; the in-person effects appear at the shock and persist, and the online count reduction is immediate. The intensive margin is omitted in dynamic form for the reason given above.




Figure E1: Dynamic treatment effects: transaction counts (PPML) by channel

Notes: Monthly PPML event-study coefficients ($\exp(\beta) - 1$, in percent) with individual and city $\times$ month fixed effects on the 434K matched sample, estimated over December 2018 through December 2021 and normalized to December 2019. Shaded bands are 95% confidence intervals with region-clustered standard errors (66 regions). The dashed line marks January 2020.




Figure E2: Dynamic treatment effects: probability of any purchase by channel

Notes: Same design as Figure E1 with the binary indicator $1[Y_{it} > 0]$ as the outcome (linear probability model); coefficients in percentage points.

F Macroeconomic Implications: City-Level Retail Sales

This appendix reports the city-level analysis of fintech credit and aggregate retail sales referenced in Section 5. Results are reported here as supplementary evidence; the identification rests on a city-month panel with limited time variation and should be interpreted with caution.

We study the macroeconomic impact of fintech credit on consumption at the city level (Table F2).

Considering the significant impact fintech credit has on consumption, along with the extensive user base, it is reasonable to expect that the utilization of fintech credit could have an impact on macroeconomic variables.

We utilized our initial random sample of nearly one million fintech credit consumers to examine the relationship between fintech credit activities and retail sales at the city level. First, we calculate the monthly *per capita fintech credit payment* for consumers in each city, which we use as a monthly indicator of fintech credit usage intensity among cities. Then we get the city-level retail sales data as reported by municipal governments²⁰ and calculate the *retail per capita* based on the residential population. Table F1 reports summary statistics for the resulting city-quarter panel. We then apply the following specification with the city-month level panel data:

$$y_{jt}^{\text{Retail}} = \gamma \cdot y_{jt}^{\text{fintech credit}} + \delta_t + \lambda_j + \epsilon_{jt} \quad (\text{F1})$$

where λ_j is the city fixed effect, capturing the city-specific fundamentals associated with retail sales, and δ_t is the fixed effect. Due to the limitations of the observations and the relatively longer time unit in this analysis, we do not distinguish the impacts of the pandemic over different cities besides the time-fixed effect. In addition, we assume the effect on social retail is time-invariant, which might be violated, particularly when the acceptance of fintech credit in different cities is changing at different paces over time.

Table F2 shows the fintech credit payment is significant in explaining the variations in retail sales at the city level for both monthly and quarterly available data. A one percent increase in consumers' fintech credit payment is associated with a 0.35 percent increase in retail sales per capita in the half-year period following the shock, and a 0.17 to 0.26 percent increase in the long term.

In column (2), we add the digital payment amount in the regression. It is not significant and does not weaken the coefficient on fintech credit, suggesting that the association reflects fintech credit usage rather than digital payment activity more broadly. However, we are not suggest-

²⁰The data is from the CEIC database. Between 2019 and 2021, 291 cities have reported the retail sales data. Of these, 137 cities reported monthly while the other cities reported quarterly. The sales in January and February are combined due to the Spring Festival holiday, and we apply a similar process to the consumer data.

ing that digital payment adoption has no effect on consumption. Recently, [Agarwal, Ghosh, et al. \(2024\)](#) found an increase in household consumption by about 2.38% for an additional 10 percentage points in prior cash dependence when adopting mobile payments. [Jack and Suri \(2014\)](#) also find an effect on consumption smoothing of mobile money adoption. In our setting, both control and treatment consumers have intensive digital payment usage, and the payment system has become highly digitalized. Therefore, building on these results, our findings emphasize the role of fintech credit in increasing consumer spending. Considering China’s digital payment advancements ([Didier et al., 2022](#)), our analysis could serve as a reference for other economies rapidly transitioning to digital payments.

Table F1: City-Level Panel: Summary Statistics

	Mean	SD	P25	Median	P75
Retail sales per capita (yuan)	6972	4314	4042	5844	8786
Fintech credit payment per sampled user (yuan)	4472	901	3843	4383	4998
Total Alipay payment per sampled user (yuan)	11874	4309	9025	11129	13738
Fintech credit share of payment volume	0.398	0.081	0.347	0.398	0.446
Residential population (millions)	4.80	3.64	2.58	4.03	5.84
City-quarter observations	1,834				
Cities	261				
Period	2020Q1–2021Q4				

Note: Summary statistics for the city-quarter estimation panel underlying Table F2, columns (1)–(2). Retail sales per capita divides reported retail sales of consumer goods by residential population; cities reporting monthly are aggregated to quarters, with January and February combined because of the Spring Festival. Fintech credit and total Alipay payments per sampled user are quarterly averages over the roughly one million fintech credit users drawn for this analysis, aggregated by city of residence. The fintech credit share is the ratio of fintech credit payment to total Alipay payment volume. Sources: CEIC city retail releases; Alipay administrative records; China City Statistical Yearbook population data.

Table F2: Fintech Credit and City-Level Retail Sales

	Dependent variable: IHS(<i>Retail per capita</i>)			
	2020-21		2020H1	
	(1)	(2)	(3)	(4)
IHS($y_{Fintech\ Credit}$)	0.18*** (0.04)	0.17*** (0.05)	0.36*** (0.10)	0.35*** (0.11)
IHS($y_{Total\ Payment}$)		0.00 (0.03)		0.01 (0.01)
City FE	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
R ² Adj.	0.972	0.972	0.983	0.983
Num. Obs.	1834	1834	405	405