

Demographics and Technology Diffusion: Evidence from Mobile Payments*

Nicolas Crouzet[†] Pulak Ghosh[‡] Apoorv Gupta[§] Filippo Mezzanotti[¶]

Abstract

This paper leverages insights from the adoption of mobile payments in India to argue that the age composition of the population can impact the diffusion of new technologies. Using data from a leading bank, we find that younger adults tend to prefer mobile payments over traditional cards. More broadly, age outweighs other observable consumer characteristics in explaining variation in the use of mobile payments across consumers. In a simple model of technology adoption, these age-driven differences in attitudes toward technology create stronger adoption incentives for businesses that are more likely to face younger consumers. We validate this prediction using store-level data on mobile payment adoption by merchants. Using both a difference-in-difference estimator exploiting actual age differences across districts and an instrumental-variable strategy using historical determinants of fertility to generate variation in age distribution across areas, we show that the adoption of the technology by businesses was higher in regions with younger population. These findings suggest that aging can pose an obstacle to the diffusion of financial innovation.

Keywords: Payments, Technology Diffusion, Fintech, Demographic Structure, UPI.

JEL Classification: O33, G23, L86, J11.

* *This version: December, 2024.* We thank Ayush Sinha for excellent research support. We also gratefully acknowledge financial support from the Financial Institutions and Markets Research Center, Kellogg School of Management. We thank participants at the UK Payment System Regulator, Clemson University, Norwegian Business School, Copenhagen Business School, Workshop on Entrepreneurial Finance and Innovation (WEFI), TPRI at Boston University, IESE Business School, Universitat Pompeu Fabra, NYU Stern - IIMC India Conference, and IPA-GPRL 2024 Research Gathering, and our discussants Sergey Sarkisyan and Boris Vallee for helpful comments and discussions. Lastly, we thank the staff of the two organizations that shared data with us for the paper for their help in the data extraction phase as well as clarifications later. The data is shared solely for the purpose of academic research. The companies did not play a role in drawing inferences, and the views expressed are solely of the authors.

[†] Northwestern University : n-crouzet@kellogg.northwestern.edu

[‡] Indian Institute of Management : pulak.ghosh@iimb.ac.in

[§] Dartmouth College : apoorv.gupta@dartmouth.edu

[¶] Northwestern University : filippo.mezzanotti@kellogg.northwestern.edu

1 Introduction

The progressive aging of the population in developed economies has recently spurred renewed research on the economic consequences of large demographic shifts. This research has shown that population aging can impact the rate or direction of innovation and, ultimately, productivity growth. Among others, [Derrien et al. \(2023\)](#) shows that young workers are often key drivers of innovation within firms, while [Acemoglu and Restrepo \(2022\)](#) and [Abeliansky and Prettnner \(2023\)](#) argue that a shrinking working-age population should spur innovation in labor-saving technologies.¹

In this paper, we study a complementary channel through which demographics and aging may impact productivity growth: the rate of *diffusion* of new technologies, as opposed to the rate of innovation itself. A large literature has argued that the adoption of new technologies is a key component of the link from innovation to growth, but also one that is subject to a number of frictions, ranging from information, to coordination, to financial frictions ([Hall and Khan 2003](#)). We provide evidence that, in the case of consumer-facing technologies, heterogeneous preferences across demographic groups, and particularly across age cohorts, play a central role in shaping diffusion rates. These effects are both direct and indirect: age accounts for a substantial part of the variation in consumers’ propensity to use technology; and heterogeneous propensities across age groups shape business decisions around the adoption of new technologies.

The context of our analysis is the diffusion of mobile payment technologies in India. We define mobile payment technologies as electronic systems allowing consumers to settle transactions using a phone or other digital device. Among electronic payment technologies, mobile payment is the main alternative to traditional bank-issued credit or debit cards. Since 2016, the rapid diffusion of mobile payment technology has dramatically altered the payment landscape in India. Prior to 2016, India’s electronic payments were predominantly facilitated by cards, similar to many developed countries. However, the Demonetization gave momentum to mobile payment options. While the initial surge was driven by the adoption of mobile wallets ([Chodorow-Reich et al. 2019](#); [Crouzet et al. 2023](#)), the Unified Payments Interface (UPI) has been the main driver of the continued diffusion of this technology in more recent years.²

Overall, the speed of diffusion of mobile payments in India stands out: between 2016 and 2020, mobile payment technologies essentially replaced cards as the main mean of electronic payments, with their share in total electronic payments increasing from less than 10% to approximately 80%, as illustrated in [Figure 1](#). Given the potential effects of mobile payment usage on financial inclusion and economic activity ([Yermack 2018](#); [Das et al. 2022](#); [Dubey and Purnanandam 2023](#); [Alok et al. 2024](#)), understanding the mechanisms behind this transition is an important question, with potential

¹Other work highlighting the link between aging and innovation, both theoretically and empirically, include [Ludwig et al. \(2012\)](#), [Hashimoto and Tabata \(2016\)](#), [Costinot et al. \(2019\)](#), [Cheng and Weinberg \(2024\)](#), and [Aksoy et al. \(2019\)](#). Relatedly, [Lewis \(2011\)](#) and [Anelli et al. \(2019\)](#) study the impact of immigration on technology choice in the manufacturing sector and on innovation, respectively. Aging populations have numerous other economic consequences beyond innovation, including labor market shortages, increased pressure on pension systems, and higher healthcare costs, potentially slowing economic growth; see [Bloom et al. \(2003\)](#) and [Gordon \(2017\)](#) for overviews.

²Section 2 explains the distinctions between mobile wallets and the UPI.

relevance to other environments beyond India.

Our analysis proceeds in three steps. First, we show that, empirically, the propensity to use mobile payment technology is strongly (and negatively) related to age, even after controlling for other potential observable determinants of technology choice. Second, we develop a simple model of technology adoption by businesses, where, consistent with the data, consumers of different ages value access to mobile payment technologies differently. Third, we test empirically the main implication of the model, namely that businesses are more likely to adopt mobile payments if they operate in markets where their potential customer base is younger. Our test uses merchant level technology adoption data from a leading Indian fintech provider of payment services, and leverages the introduction of new payment modalities in 2019. Our evidence strongly supports the view that the technology we study diffused faster in districts where the customer base was younger, and highlights the importance of demographics and consumer preferences in shaping technology diffusion.

The first step in our analysis is to document the strong empirical relationship between customer age and the propensity to use mobile payments. We use a dataset comprising approximately 200,000 customers from one of India’s largest banks. The data includes comprehensive bank account activity and demographic information for a subset of customers.³ In particular, it allows us to measure the proportion of electronic payments made using mobile technologies. We establish two main stylized facts. First, we use a Shapley R-squared decomposition to quantify the contribution of age to cross-sectional variation in the mobile payment share, relative to other economic or demographic covariates. We show that age is a primary contributor to this variation: it accounts for about 38% of total cross-sectional variance in mobile payment use, much more than wealth (7%) or occupation (5%).⁴ Second, we show that younger customers strongly favor mobile payment relative to older consumers. The relationship with age is largely monotonic, robust to controlling for a host of factors, including occupation, marital status, assets, location, or even access to credit cards. It is also quantitatively large: the share of mobile payments is half as large in the oldest age bracket (60 and older) than in the youngest one (30 and younger). Using an alternative data set, we also show that the gap in the use mobile payments between young and old appeared in the data already in the early phase of mobile payments (e.g., 2016), and has persisted at least until 2022, when our data ends.⁵

In the second step, we develop a simple structural model to work out the implications of these differences in propensity to use technology across age demographics for the adoption decisions of businesses. In the model, businesses must decide whether to invest in a new technology to process sales, which we interpret as mobile payments. They face customers of two potential types, young

³The dataset we used is described in detail by [Agarwal et al. \(2022\)](#). Although our sample generally represents individuals who are wealthier than the average Indian citizen, the age distribution within our sample closely aligns with the national demographic distribution.

⁴The other primary contributor is geographic location (i.e., six-digit pincode), which accounts for approximately the same amount of variation in the payments share as age.

⁵While quantifying the mechanisms explaining this relationship is outside the scope of the paper, [Section 2](#) provides some discussion of this issue.

and old. To clarify the analysis, we assume that these two groups only differ in one dimension: young consumers' preference are sensitive to the technology choice of the business with whom they interact for their purchases, while old consumers are not. We represent this difference as taste shifters, and assume businesses can make investment to induce changes in these taste shifters, which we interpret as the technology adoption choice. In equilibrium, each consumer only makes purchases from one business, but each business sells to a group of consumers of different ages.

We show that, the lower the typical age of the consumer that the business expects to service, the higher the rate of adoption of the technology by businesses overall. Intuitively, technology improvements increase the likelihood that the business will attract young customers.⁶ Additionally, the model makes the prediction that factors affecting the adoption cost of the technology (including marketing campaigns or financial incentives) will have a lower impact on the take-up rate of the technology in environments where the customer base is older. More generally, in service-oriented sectors, higher age among consumers could either directly slow down technology adoption, or amplify other factors that contribute to slow diffusion, such as adoption externalities.

Finally, in the third step of our analysis, we provide direct empirical evidence consistent with business technology adoption decisions being influenced by the demographics of their customer base. We study the introduction of QR code-enabled terminals by a prominent fintech company in India in 2019. This company offers payment processing services to merchants, providing them in particular with point-of-sale (POS) machines. Until 2019, the functionality of these terminals was limited to traditional card payments. However, in May 2019, the Company expanded its offerings to include terminals capable of processing payments via QR codes, thus accommodating mobile payment applications. This shift allows us to assess whether — consistent with the model — merchants' propensity to adopt mobile payments is influenced by the demographic structure of potential customers. In particular, we study how the adoption of our company's services change after the May 2019 policy in relationship with the share of young adults in the area, which we define as the share of the population less than 30 years of age.

In our baseline results, we find that, on average, a one-standard-deviation higher share of young adults is associated with a 25% higher adoption response to the introduction of the QR code option. Importantly, this increase does not materialize until two months after the announcement of the new option, and is not explained by differential adoption patterns before the announcement. Additionally, our results are generally robust to changes in the specification and definition of the treatment variables. Our interpretation of these results is that merchants face a stronger incentive to adopt mobile payments in districts with younger demographics because of intrinsic preferences of young consumers for the technology.

We also provide three sets of results that speak to the potential criticisms which might be directed at this interpretation. One such criticism is that the increased adoption observed in younger districts simply reflect a correlation between age and other demographic or economic

⁶We assume a homogeneous price elasticity between young and old consumers, implying that the average markup charged by businesses is independent of the demographic composition of their customer base. An appendix extension shows that our main results survive so long as young consumers are more price-elastic than old ones.

characteristics of the district which themselves influence business adoption decisions. For example, if younger districts are more educated or wealthier, then education or wealth — not age — may be the key traits driving higher adoption rates. Our data do show that younger districts differ from others along a number of dimensions, but notably, these districts are, on average, less affluent and less educated. But most importantly, we show that controlling for these correlated demographic and economic characteristics does not alter our baseline results, and, if anything, strengthens them.

Moreover, we find similar results when we instrument our main treatment variable using historical determinants of fertility. To be precise, we leverage the simple observation that the presence of a skewed sex ratio in a region, should predict — all else equal — lower birthrates going forward (Guilmoto 2012; Dyson 2012; Angrist 2000). We then use a quadratic function of the sex ratio in 1991 to instrument the share of young adults two decades later.⁷ This approach allows us to isolate variation in the youth share that is driven by historical demographic features and should be orthogonal to recent local migration trends. Using this approach, our main results are confirmed both qualitatively and quantitatively.

To support our interpretation further, we propose a novel approach that utilizes the distribution of universities within a district. The premise is that businesses in university areas inherently serve a younger clientele (i.e., students). Yet, these businesses should resemble others within the same district but different neighborhoods. By analyzing the presence of universities at the pincode level—the most detailed location definition in our data—and incorporating district-by-month fixed effects, we demonstrate that adoption rates post-May 2019 indeed rose more significantly in university pincodes compared to other pincodes within the same district.

The main advantage of this approach is that the district-by-time fixed effects help to exclude any time-varying district-level confounders, such as variations in business owner composition or other supply-side factors. Another benefit is that it clearly identifies the consumer group driving the demand for the technology: university students. This specificity allows us to predict which types of businesses will be most influenced by this dynamic. Consistently, our findings are primarily driven by businesses that cater to university students locally. Conversely, we observe no significant effects on merchants less likely to serve students, effectively serving as a placebo group.⁸

Thus overall, we show that age is a key determinant of the use of mobile payment technology in India, and that, consistent with a simple model of technology adoption, business adoption decisions reflect age differences in their potential customer base. More generally, our evidence supports the view that demographics may be an important driver of the diffusion of new technologies. A key advantage of our approach, relative to cross-country comparisons, is that we side-step the

⁷The quadratic function allows us to flexibly exploit the effects of skewed sex ratio (Hesketh and Xing 2006; Hesketh and Min 2012), without imposing any structural assumption about the optimal level of sex ratio or the slope of this relationships.

⁸For instance, this approach allows us to assuage the concerns around the importance of differences in managers' demographics in explaining our results. First, while areas with universities should be characterized by a younger clientele, the makeup of the business owners should not be different. Second, if still differences in the demographic composition of managers should be the reason for our findings, the impact of this alternative mechanism should be picked up by our placebo test.

challenges created by differences in technologies and institutional background across countries. The competitive and regulatory framework across the Indian districts we compare is relatively homogeneous, while the technology we study is exactly the same, improving our ability to isolate the effects of demographics relative to other factors.

Section 2 reports the stylized facts on the relationship between age and the propensity to use mobile payments. Section 3 outlines a model connecting this propensity to business adoption decisions. Section 4 describes the evidence on the effect of population age on merchants' decision to adopt mobile payments. Section 5 concludes by discussing the broader implications of our evidence.

Contribution to the literature Our findings contribute to three bodies of literature. First, they relate to the literature on the adoption and diffusion of new technology. While previous studies have documented that age differences can affect a customer's decision to adopt a new product (Klee 2008; Wang and Wolman 2016), our key contribution is to extend this analysis and demonstrate that these differences can have a significant economic impact on merchants' decisions to adopt. In other words, our results underscore how heterogeneity in consumer technology preferences can influence firms' adoption decisions.

Existing work has predominantly focused on the agricultural or farming sectors, where the preferences of end consumers regarding the technology used in production are less relevant, provided they do not affect the final product's quality or price.⁹ By contrast, consumer preferences may be more important determinants of adoption decisions in the service sector, since the technology used to deliver services can be an integral part of value creation, making consumer preferences crucial. Importantly, our results highlight that differences in preferences will not only affect adoption due to demand differences among final consumers, but will also reduce technology adoption on the business side. An implication is that differences in diffusion rates could derive from differences in consumer preferences, including those driven by demographic characteristics, which are the focus of our analysis. Thus our evidence adds more broadly to the literature on why new technology diffuses slowly, even when financial, regulatory, or informational hurdles are not obvious (Hall and Khan 2003; Comin and Hobijn 2010; Foster and Rosenzweig 2010; Manuelli and Seshadri 2014).

Second, our research contributes to the fintech literature, which has seen a surge of interest in analyzing the drivers and impacts of new payment technologies. A significant body of recent research has focused on understanding the expansion of various payment methods, including traditional cards (Higgins 2023; Aggarwal et al. 2023), crypto (Hu et al. 2019), mobile wallets (Chodorow-Reich et al. 2019; Crouzet et al. 2023; Vallee et al. 2024), and instant payment systems like UPI in India (Dubey and Purnanandam 2023; Alok et al. 2024) and Pix in Brazil (Sarkisyan 2023).¹⁰

⁹For instance, Atkin et al. (2017) studies the role of organizational constraints in the manufacturing of soccer balls; work by Munshi (2004), Conley and Udry (2010), and Gupta et al. (2022) examine from these perspectives the role of information frictions and learning in agriculture. Our findings also relate to Goehring et al. (2023), which studies the role of career concerns in technology adoption.

¹⁰This paper also complements the literature studying the real impact of electronic payments. For instance, several papers have examined the impact of digital payments on households' behavior (Jack and Suri 2014; Suri and Jack 2016; Bachas et al. 2021; Aggarwal et al. 2024; Bian et al. 2023) and businesses (Agarwal et al. 2019).

Despite the wealth of insights, a common characteristic of these studies is that they focus on a specific electronic payment method (relative to cash). Our study diverges from most of the previous literature by examining the decision-making process between different electronic payment options. Our results suggest that the simultaneous presence of multiple technologies (i.e., multi-homing) could partially arise from heterogeneous consumer preferences for distinct products.¹¹

Finally, our paper is connected to work on the productivity implications of large demographic transitions (Feyrer 2007, 2008; Acemoglu and Restrepo 2017; Maestas et al. 2023; Acemoglu and Restrepo 2022). A related literature has connected aging to declining rates of entrepreneurship and firm entry (Liang et al. 2018; Peters and Walsh 2019; Azoulay et al. 2020; Bornstein 2021). Our paper complements this work by providing empirical support for a new channel through which aging could affect productivity growth, distinct from entrepreneurial innovation: the diffusion of new technologies to businesses (both incumbents and new entrants).

2 Age and the propensity to use mobile payments

This section provides stylized facts on the relationship between consumer age and the propensity to use mobile payments. Our data focus on the Indian market, so we start with a brief institutional background on mobile payments in India.

2.1 Institutional background: mobile payments in India

The Indian mobile payment landscape offers a captivating example of rapid adoption of new financial technologies within a short timeframe. This section reviews these recent changes, highlighting the difference between mobile payment technologies and the other form of electronic payment technology available in India: traditional card-based transactions.

Mobile payment modalities In the Indian context, mobile payment can refer to two key technologies: mobile wallets and the Unified Payments Interface (UPI). Mobile wallets function as preloaded payment technologies, allowing users to deposit funds in their digital wallets for use in future transactions. Similarly, businesses can utilize digital wallets to receive payments. The contents of the wallets can then be transferred to the traditional bank deposit accounts of consumers and businesses. These services, often free for consumers, have attracted numerous providers competing based on security, convenience, and integration with traditional payment methods. Initially introduced in the early 2010s with platforms like Paytm and MobiKwik, their popularity surged

¹¹Our paper also relates to Jiang et al. (2024), which studies the disparate impact of digital banking among U.S. consumers, showing that older and poorer consumers are, on average, net losers in the shift from branch banking to mobile. While our papers differ in many, important dimension, a fundamental distinction is that Jiang et al. (2024) take a supply-side approach, focusing on how banks' strategic decisions (i.e., branch closure) drive the differential impact of digital transformation across groups. In contrast, our paper adopts a demand-side perspective, emphasizing how consumers' inherent preferences and characteristics, particularly age, shape technology adoption among the businesses in the market.

after India’s 2016 demonetization, with mobile payment volumes nearly tripling from April 2016 to April 2017 (Chodorow-Reich et al. 2019; Crouzet et al. 2023).

Mobile payments can also refer to the UPI. Introduced by the National Payments Corporation of India (NPCI) in 2016, the UPI facilitates immediate, real-time bank-to-bank transfers, enabling transactions via a mobile interface without requiring physical cards or certificates (Dubey and Purnanandam 2023). Managed by the NPCI, the UPI is accessible through various popular apps, including those offering mobile wallet services. Like mobile wallets, UPI services are free for consumers. Competing apps distinguish themselves through additional services or a differentiated user experience. The UPI offers two primary advantages over mobile wallets. First, it provides direct connectivity to a funding source (e.g., a bank account), eliminating the need to preload funds into a digital wallet. Second, the UPI guarantees interoperability across different banks and financial service companies.¹² While the UPI was formally introduced in 2016, UPI transactions remained small, compared to mobile wallet transactions, until the end of 2017.¹³ However, the UPI’s growth trajectory surpassed that of mobile wallets post-2017, reaching approximately 80% of mobile transactions by the end of 2021.¹⁴

Mobile payments and traditional card-based payments In addition to mobile payments, Indian households have long had access to traditional card-based electronic payment methods. Much like in the United States, Indian consumers enjoy a range of options including debit, credit, and prepaid cards. The Indian market is served by major international card companies, reflecting a level of accessibility comparable to that seen in the United States.¹⁵ Since the bulk of our analysis is concerned with comparing adoption of mobile payment technology with traditional card-based electronic payment methods, it is important to clarify the differences between these technologies. We highlight three main differences.

First, mobile payment options generally involve lower adoption costs for consumers. Typically, there are no financial expenses associated with opening a mobile wallet or registering with UPI. Moreover, the non-monetary costs involved in setting up these accounts are often less burdensome than those required for obtaining a card.

Second, the expenses borne by businesses in accepting mobile payments are typically lower compared to card transactions. For a business, the fees associated with using the UPI mobile payments may vary depending on the payment company handling the transaction, but they are generally lower than those associated with card payments.

The third significant distinction between mobile payments and cards pertains to the transaction

¹²In other words, a traditional mobile wallet managed by a fintech company A can only send money to other wallets managed by A. Instead, with UPI, you can pay any UPI holders, irrespective of the application that the firm uses to manage the UPI account.

¹³<https://www.npci.org.in/what-we-do/upi/product-statistics>

¹⁴The Reserve Bank of India (RBI) provides aggregate statistics about payment, allowing to separately measure the amount of UPI and mobile wallets. More discussion on the data is provided when we present Figure 1.

¹⁵For instance, the major card providers (i.e., Visa, AMEX, and Mastercard) all operate in the Indian market. One notable difference with the United States market is that the Indian government has entered indirectly the offering of card services through Rupay.

process itself. As the term suggests, mobile transactions are executed using an app on a phone or similar digital device.¹⁶ In consumer-to-business transactions, QR code technology is the primary payment method, allowing consumers to swiftly complete purchases by scanning a QR code provided by the merchant, and facilitating rapid and contactless payments.¹⁷ Additionally, the digital interfaces of applications hosting the UPI or mobile can offer a customized consumer experience, with additional options to monitor payments made or transfers received in real-time, for instance.

The expansion of mobile payments Aggregate data from the Reserve Bank of India (RBI) underscores the remarkable surge in mobile payments that occurred from 2016 onward (Figure 1). Prior to 2016, India’s electronic payment landscape was largely dominated by card-based transactions. However, this landscape underwent a significant transformation following the Demonetization at the end of 2016. Not only did this event spur a general increase in electronic payments, but it also notably bolstered mobile payments, primarily through mobile wallets. The momentum towards mobile payment dominance persisted beyond 2017, with UPI transactions gradually capturing a larger share of mobile payment volumes. By 2019, mobile payments equaled the volume of card transactions and have since continued to grow at a rapid pace. As of the end of 2021, mobile payments represented the predominant form of electronic payment in the Indian market.

The Indian transition of electronic payments from card-based technologies to mobile technologies is particularly striking when contrasted with the recent evolution of electronic payments in many developed countries, including the United States. Recent market research shows that in 2023, Apple Pay, the most popular mobile payment option in the US, only accounted for 3.1% of all in-store purchases in the United States by volume, indicating a comparatively low rate of adoption of the technology by consumers.¹⁸ Additionally, it is crucial to recognize that most mobile payment options in the United States are still fundamentally linked to credit cards, and therefore represent a smaller step in innovation than mobile payments in India.¹⁹

2.2 Consumer age and the propensity to use mobile payments

Many economic, technological, and institutional factors could explain the recent surge in mobile payments in India. Our focus in this paper is on the role of age. Our key premise is that young

¹⁶In theory, both mobile wallets and UPI have options that do not require a smartphone, allowing payment validation through a phone call or text, but this option appears relatively uncommon (albeit exact statistics are hard to find).

¹⁷While credit cards theoretically can be integrated into a digital interface for use via QR code scanning, akin to how ApplePay operates in the US, this digital card option appears relatively rare within our context. For instance, in the dataset provided by our fintech company used later for the analyses, we found that a small percentage (3% of volume) of QR code transactions were conducted using cards in 2019.

¹⁸See <https://capitaloneshopping.com/research/apple-pay-statistics>. A 2021 survey by PYMNTS confirms this qualitative fact: this survey adopts a wider definition of mobile wallet (i.e., not only Apple Pay but also other providers) and finds that only about 10% of US respondents had recently utilized this payment option; see <https://www.pymnts.com/apple-pay-tracker/2021/7-years-later-6pct-people-with-iphones-in-us-use-apple-pay-in-store/>.

¹⁹In other words, services like Apple Pay build on the the pre-existing card network, rather than replacing it. If anything, this feature should make scaling easier.

consumers tend to be more predisposed to use mobile payment technologies. In a country with a younger population, this predisposition could not only directly generate more mobile payment usage, but also, indirectly, encourage greater adoption among businesses. The remainder of this section presents evidence consistent with our argument’s foundation: namely, that younger consumers exhibit a significantly higher propensity to use mobile payments, and that the association between age and mobile payments usage reflects intrinsic preferences, as opposed to other demographic, economic, or geographical factors potentially influencing mobile payments usage.

1. *Data sources*

Our primary dataset comes from one of the top four banks in India, encompassing approximately a sample of 200,000 customers. This bank operates an extensive network of over 18,000 branches and ATMs, offering a comprehensive suite of financial products and services.²⁰ The dataset used in this study contains transactions from January and February 2020 and provides insight into the usage of traditional cards versus mobile payments, with the latter measured solely through UPI transactions.²¹ Additionally, the data provides basic demographic information about the clients. As our focus is on understanding how age effects impact payment preferences, we compare the age distribution of the dataset with the national demographic profile of household heads, as reported in the National Family Health Survey (NFHS) from 2019-2021.²² As illustrated in panel (a) of Figure A-1, the age profiles of bank account owners and household heads closely align, with a minor under-representation of individuals aged 60 to 65 offset by a higher presence of middle-aged individuals (30-50 years). Lastly, we note that our data are mechanically skewed toward wealthier household, since households must maintain a bank deposit account to be in our sample (panel b of Figure A-1). However, the data has a relatively broad coverage of wealth levels, allowing us to disentangle the effects of age from those of wealth.²³

2. *Results*

Age as a source of variation in mobile payment usage We start by documenting the degree to which age accounts variations in payment preferences when contrasted with other factors, such as gender, occupation, marital status, wealth, or geographical location. To do this, we employ a Shapley R-squared decomposition method (Huettner and Sunder 2012; Israeli 2007).²⁴

²⁰To maintain confidentiality, we refrain from disclosing the bank’s identity, although its data has been utilized in other academic studies, such as Agarwal et al. (2022).

²¹Card payments include transactions made with both debit and credit cards, while mobile payments are determined by UPI transactions.

²²We choose to compare our data’s age distribution with that of household heads as this characteristic is more likely to be the one comparable to our measure. In fact, most households have only one account, typically under the head’s identity. Throughout our analysis, we only consider bank customers aged between 18 and 65.

²³Appendix A.2.1 contains a more detailed discussion of the comparison of the age and wealth distribution of our data with nationally representative samples of households.

²⁴Several applied papers have used this method in recent years, including Biasi and Ma (2022) and Mezzanotti and Simcoe (2023). Drawing on the concept of the Shapley value in cooperative game theory, this method calculates the average marginal contribution of a predictor (in our case, age) to the total R-squared of regressions including

The findings are presented in Table 1. These results highlight age as the primary economic or demographic factor explaining the largest share of variance in payment methods.²⁵ The precise contribution of age to the variation in mobile payment share of depends on the other factors included in the decomposition. Nevertheless, we can use the most conservative estimates, obtained by including all factors simultaneously, as a benchmark. In this scenario, age accounts for approximately 38% of the explained variance, making it the most significant factor alongside location (i.e., six-digit pincode), which explains roughly 42% of the variation. Marital status follows as the next significant factor (8%), trailed by wealth (7%) and occupation (5%).²⁶ The depositor’s gender proves to be essentially inconsequential. Thus age emerges as a key characteristic accounting for the cross-sectional variation in payment preferences between cards and mobile payment.

Mobile vs. Card across the Age Distribution We now analyze the relationship between payment preferences and the age of the account holder using the data. To ensure clarity, we start by documenting how the proportion of electronic payments made via mobile varies across different age groups without controlling for additional covariates. Specifically, Figure 2, panel (a), reports a non-parametric scatter plot of the relationship between the share of mobile payment amounts and age. We observe a negative, monotonic, and approximately linear relationship between age and mobile payment usage: older individuals consistently utilize mobile payments less frequently than cards. These differences are substantial, with consumers in the oldest category conducting approximately 25% of their electronic payments using mobile, compared to 55% for younger consumers.²⁷

Next, we introduce individual-level controls. The objective is to disentangle the effect of age from other observable characteristics that may influence electronic payment preferences and could be correlated with age. In Figure 2, panel (b), we incorporate demographic controls for gender, marital status, and occupation. These controls are applied by residualizing them against both the proportion of payments made via mobile and age, and then plotting the residuals against age. This adjustment has minimal impact on the observed relationship: indeed, the coefficient in the linear fit of the relationship (reported in the figure) remains virtually unchanged from panel (a). In Figure 2, panel (c), we further introduce controls for the wealth of the bank customer to mitigate the possibility that age-related differences are merely reflections of wealth disparities across cohorts. Once again, the inclusion of this control has a relatively modest impact.²⁸

all possible subsets of predictors, thus offering a breakdown of the total R-squared among all combinations of the predictors considered. In our case, the additional predictors beyond age include gender, marital status, occupation, wealth (proxied by total deposits), and location, defined by a (6-digit) pincode.

²⁵Age is defined non-parametrically using age groups, with 48 dummies classifying all ages between 18 and 65.

²⁶The impact of marital status is intriguing and may indicate differences in the number of individuals using the bank account between married and single individuals.

²⁷In Appendix Figure A-2, we replicate the same analysis using age groups (i.e., 18-25, and then at 5-year intervals) and present the results with confidence intervals relative to zero. This confirms that mobile payment usage significantly differs from the youngest group for every age group, with each subsequent age group exhibiting lower mobile usage than the preceding one.

²⁸We total account balances (including savings in fixed deposits, mutual funds investments, public provident funds accounts, recurring deposits accounts, and savings accounts) held by the customers with the bank as an empirical proxy for their wealth. We then control for wealth by creating 20 equal bins each month and then using fixed effects

We then introduce controls for location. Different age groups may reside in distinct parts of the country or in different neighborhoods within the same districts. For instance, the younger population may locate in areas where stores are less inclined to accept credit cards, potentially increasing their reliance on mobile payments. In that case our results would reflect lack of access to credit card payments, as opposed to a preference for mobile payments. To address this concern, Figure 2, panel (d), replicates the previous analysis but includes controls for pincode-by-wealth group fixed effects, alongside standard demographic controls.²⁹ Although the magnitude of the relationship between age and mobile payment usage is somewhat diminished, consistent evidence of a significant relationship between age and mobile payment usage remains compelling.³⁰

A final concern is that age is a proxy for differences in the ability to obtain a card across different age groups. Older individuals might be more likely to be approved for debit or credit cards, potentially underpinning the observed relationship.³¹ To address this issue, Figure 2, panel (e), conducts a similar analysis as before — incorporating individual controls and pincode-by-wealth fixed effects — but focuses only on customers who possessed cards during the analyzed period. Even after conditioning on ownership of a card, we find that younger consumers consistently allocate a significantly higher proportion of their expenditures to mobile payments. Specifically, the share of mobile payments is approximately 30% higher for the youngest cohort than for the oldest one. The linear fit of this relationship remains quantitatively identical to the one estimated in panel (a).³²

Discussion After highlighting the surge of mobile payments in India in recent years, this section shows that consumers of different ages exhibit distinct propensities to use mobile payments. The relationship between age and mobile payment usage is both economically significant and broadly monotonic. This relationship persists even after controlling for differences in occupation, wealth, geographic location, and electronic payment card ownership.

We interpret this evidence as suggesting that younger consumers in India have a stronger preference for mobile payments than their older counterparts. A natural question is what mechanisms may explain this difference. On the one hand, the behavior could reflect cohort-specific preferences for certain features of new technologies. For instance, younger cohorts may naturally be more inclined to adopt and use digital technologies (Prensky 2001a,b). On the other hand, the preference for mobile payments could stem from differences in life experiences across cohorts. Older adults, having been more exposed to older payment technologies (i.e., cards), may have developed habits that create a preference for continuing with these options, even when newer ones are available

for each of the 20 bins.

²⁹Pincodes are at the 6-digit level, so the fixed effects are expected to significantly mitigate variation in business types encountered by individuals.

³⁰Additionally, with the full set of controls, we repeat the analysis using constructed age bins rather than equal-sized bins and find similar results.

³¹It is important to note that differences in card ownership among cohorts may also stem from varying preferences. For instance, if a young person strongly prefers mobile payments, they may choose not to apply for a credit card. This suggests that the test conducted here may, in part, underestimate the role of preferences, as defined in this study.

³²Due to the smaller sample size and the large number of controls, the relationship between payment and age is slightly noisier, particularly at higher age levels.

(Dyanan 2000).

Although this paper does not aim to fully explain the source of heterogeneity, a few clarifications are in order. In general, the importance of identifying the exact mechanism arises primarily when making predictions about the long-run dynamics and is less crucial for explaining adoption patterns in the cross-section or during the early stages of the technology’s life. In fact, the main distinctive difference between the two explanations is in the time-series: in particular, differences in behavior due to habits are likely to diminish over time, leading the gap between the two demographic groups to shrink in the future. In a static framework, this feature is clearly not critical, and more broadly, it becomes important only as sufficient time passes for the role of habits to wane.

In this context, we introduce two new tests that aim to better assess the extent to which habit formation may play a role in explaining our findings. While the data suggests that habit formation may partly explain the choice between payment forms, we also find that differences in habit do not dissipate quickly: the gap between the young and old has persisted for several years and shows no significant reversal in recent years. To make these claims, we use a panel data from a from the same bank.³³ While this sample focuses on a more limited set of clients in a few large districts in India, it tracks their monthly spending behavior from 2012 to 2022, and therefore allows us to explore the dynamics of the preference. We conduct two complementary exercises with this sample.

First, we provide evidence consistent with habit formation by examining whether early use of a credit card is associated with a lower likelihood of using mobile payments in the later part of the sample. In Figure A-3, we plot the month-by-month difference in the probability of using mobile payments across individuals with and without a credit card in 2015. After mobile payments gained significant traction (i.e., post-2017), we observe a lower likelihood of using mobile payments among individuals who already had credit cards. We find that early users of cards keep using mobile payment at a lower rate, at least until 2021.

Next, we examine how young clients’ preference for mobile payments has evolved over time. In the panels of Figure A-4, we plot the monthly difference in the average use of mobile payments between younger and older clients starting in 2016. Throughout the sample period, younger clients consistently used mobile payments more than their older counterparts. This gap widened during the boom period for mobile payments and stabilized around 2020 for this sample. While mobile payments grew significantly during the analysis period (Figure 1), the level of penetration remains different across cohorts. To the extent that habits may have explained in part the difference between cohorts, this evidence suggests that this mechanism does not dissipate quickly, since we find large differences in behavior six years after the boom in mobile payments.³⁴

³³While the data carries the advantage of providing transactional details over a long panel of customers, including providing details before 2016, it is nevertheless not well-suited for our main analysis for few reasons. First, the data comes only from nine biggest metropolitan Indian cities and, therefore, is not representative of less urban regions. Second, because the data keeps the panel of customers fixed beginning in 2012, we are unable to capture a large share of younger population in the year 2020 with this data. The test discussed below use the sample of customers that entered the data before 2015 and kept using the account until at least 2021.

³⁴Importantly, we find that the persistence of the age gap is stronger than the persistence in the difference in behavior found between early card-users and consumers without a card in 2015. Arguably, this comparison is a more direct test of habits. This suggests that the age-induced preference appears to persistent longer than what induced

Altogether, our evidence suggests that the preference for mobile payments may partly stem from differences in habits between groups. However, we also find that—to the extent habits contribute to the age gap—their importance does not appear to dissipate even after several years. In general, the behavioral difference between age groups increased during the early phase of the mobile payment boom and has remained stable in recent years. This stability implies that the diffusion effects discussed in the next section are unlikely to be short-lived and will likely persist over the medium run.

3 Model

In this section, we outline a simple model of the interaction between demographics and the adoption of new technologies by businesses. The model shows how differences in the age structure of the population can lead to different rates of technology adoption by businesses when consumer attitudes toward technology vary with age, as documented in Section 2.

3.1 Exposition

The model is static. There is a continuum of mass 1 of consumers indexed by $i \in [0, 1]$. All consumers have the same income, which, for simplicity, we assume to be given by the wage rate w . Additionally, there are $j = 1, \dots, J$ businesses operating in the economy. The number of businesses is fixed — there is no entry —, and in equilibrium all of them will operate. Businesses hire labor at the wage rate w for their operations. The model is set in partial equilibrium, by which we mean that we take the wage rate w as given, and do not require that labor markets clear.

1. Consumers

We assume that each consumer can only make purchases from one business, though they can choose which business to make these purchases from. The problem of each consumer is to choose a business, j , and a quantity to purchase, $c(i, j)$, to solve:

$$\begin{aligned} \max_{j \in [1, \dots, J], c(i, j)} & \quad \frac{\log(c(i, j))}{\nu - 1} + \varepsilon(i, j) \\ \text{s.t.} & \quad p(j)c(i, j) \leq w \end{aligned}$$

Here, $p(j)$ denotes the price charged by business j , and $\nu > 1$ is a structural parameter such that $\nu/(\nu - 1)$ characterizes consumers' elasticity of substitution across businesses.

For each consumer i , the $(\varepsilon(i, j))_{j \in [1, \dots, J]}$ represent taste shifters that govern relative preferences across businesses. These shocks are drawn independently across businesses. Moreover, within each business, the shocks are drawn independently across consumers.

by a direct early use of cards.

There are two types of consumers: young and old. Let \mathcal{I}_O denote the set of old consumers, and $\mathcal{I}_Y = [0, 1] \setminus \mathcal{I}_O$ denote the set of young consumers. Our first key assumption is that young and old consumers only differ in their sensitivity to technology adoption choices made by businesses. We represent this difference as follows.

Assumption 1 (Preferences for technology). *Fix a business j . For old consumers, the taste shifters with respect to products sold by j are drawn from a standard Gumbel distribution, with CDF:*

$$\mathbb{P}(\varepsilon(i, j) \leq z) = \exp\left(-\exp(-z)\right) \quad \forall z \in \mathbb{R} \quad \text{if } i \in \mathcal{I}_O.$$

Instead, for young consumers, the taste shifters are drawn from a Gumbel distribution with CDF:

$$\mathbb{P}(\varepsilon(i, j) \leq z) = \exp\left(-\exp(-(z - \ln(a(j))))\right) \quad \forall z \in \mathbb{R} \quad \text{if } i \in \mathcal{I}_Y,$$

where $a(j) \geq 1$ represents the technology adoption choice of business j .

Technology adoption $a(j)$ by each business is endogenous and will depend on consumer demographics. However, consumers take technologies $\{a(j)\}_{j=1}^J$ as given. For business j , increasing $a(j)$ leads to a first-order stochastic shift in the distribution of the taste shifters of potential customers. Assumption 1 captures the idea that young consumers are more sensitive to businesses' technology offering than old consumers. For the particular case of mobile payments, this assumption is consistent with the evidence presented in Section 2.2, which highlighted the quantitative importance of the negative relationship between age and the propensity to use the technology in India.

It is clear that each consumer will spend all their income, so that:

$$c(i, j) = \frac{p(j)}{w},$$

Thus the optimal choice of which business to purchase from is the solution to the problem:

$$j(i) = \arg \max_{j \in [1, \dots, J]} -\frac{\log(p(i, j))}{\nu - 1} + \varepsilon(i, j)$$

Standard derivations, reported in Appendix A.1, then give the following result.

Lemma 1 (Demand shares). *Suppose that business j charges price $p(j)$ and makes technology adoption choice $a(j)$. Then the probability that a young consumer $i \in \mathcal{I}_Y$ chooses business j is given by:*

$$s_Y(p(j), a(j)) = \frac{a(j)}{J} \left(\frac{p(j)}{P_Y}\right)^{-\frac{1}{\nu-1}} \in [0, 1],$$

where the price index P_Y is defined as:

$$P_Y \equiv \left(\frac{1}{J} \sum_{j=1}^J a(j)p(j)^{-\frac{1}{\nu-1}}\right)^{-(\nu-1)}.$$

The probability that an old consumer $i \in \mathcal{I}_O$ chooses business j is given by:

$$s_O(p(j)) = \frac{1}{J} \left(\frac{p(j)}{P_O} \right)^{-\frac{1}{\nu-1}} \in [0, 1],$$

where the price index P_O is defined as:

$$P_O \equiv \left(\frac{1}{J} \sum_{j=1}^J p(j)^{-\frac{1}{\nu-1}} \right)^{-(\nu-1)}.$$

2. Businesses

In what follows, we denote by θ the share of young consumers in the total population. By the law of larger numbers, each business $j \in [1, \dots, J]$ faces a demand for their product given by:

$$D(p(j), a(j)) = (1 - \theta) s_O(p(j)) \frac{w}{p(j)} + \theta s_Y(p(j), a(j)) \frac{w}{p(j)}.$$

The first term captures demand from old consumers, who are insensitive to business j 's technology adoption choice. The second term captures demand from young consumers, who are sensitive to it.

Businesses all face a common and constant unit cost of sales. For notational convenience, we assume that it is given by the wage rate, w . Additionally, businesses face a cost of implementing technology choice $a(j)$ which we model as follows.

Assumption 2 (Technology adoption costs). *Choosing technology adoption level a requires $c(a)$ units of labor, where $c : [1, +\infty) \rightarrow \mathbb{R}^+$ is a twice-differentiable, strictly increasing, and strictly convex function satisfying $c(1) = c'(1) = 0$.*

This is the second key economic assumption: adopting the technology — that is, choosing a value of $a(j)$ strictly larger than 1 — imposes costs on the business. One interpretation of these costs is that the technology may require workforce training to be deployed. Another interpretation is that businesses may be uncertain that the technology is reliable. Economically, this cost creates a force that limits the scale of technology adoption by businesses.

Profits for business j are given by:

$$\Pi(p(j), a(j)) = \left((1 - \theta) s_O(p(j)) \frac{w}{p(j)} + \theta s_Y(p(j), a(j)) \frac{w}{p(j)} \right) (p(j) - w) - wc(a(j))$$

Define the markup over sales as:

$$\mu(j) \equiv \frac{p(j)}{w},$$

and define aggregate markups for each consumer group:

$$\mathcal{M}_Y \equiv \frac{P_Y}{w} = \left(\frac{1}{J} \sum_{j=1}^J a(j) \mu(j)^{-\frac{1}{\nu-1}} \right)^{-(\nu-1)}, \quad \mathcal{M}_O \equiv \frac{P_O}{w} = \left(\frac{1}{J} \sum_{j=1}^J \mu(j)^{-\frac{1}{\nu-1}} \right)^{-(\nu-1)}. \quad (1)$$

With this notation, profits scaled by the wage rate only depend upon markup and the technology adoption choice, $\Pi(p(j), a(j))/w \equiv \pi(\mu(j), a(j))$, and can be written as:

$$\pi(\mu(j), a(j)) = \left((1-\theta) \frac{1}{J} \left(\frac{\mu(j)}{\mathcal{M}_O} \right)^{-\frac{1}{\nu-1}} + \theta \frac{a(j)}{J} \left(\frac{\mu(j)}{\mathcal{M}_Y} \right)^{-\frac{1}{\nu-1}} \right) \left(1 - \frac{1}{\mu(j)} \right) - c(a(j)).$$

Finally, we assume that businesses are small (or non-strategic), so that they do not internalize the effects of their price or technology adoption decision on the level of aggregate markups, but instead take them as given. Thus the profit maximization problem for each business is:

$$\hat{\pi}(j) = \max_{\mu(j), a(j) \geq 1} \pi(\mu(j), a(j))$$

The two following necessary first-order conditions are:

$$\mu(j) = \nu, \quad (2)$$

$$c'(a(j)) = \frac{\theta}{J} \left(\frac{\mu(j)}{\mathcal{M}_Y} \right)^{-\frac{1}{\nu-1}} \left(1 - \frac{1}{\mu(j)} \right). \quad (3)$$

The markup is set equal to the price elasticity of demand of the consumer base of each business. Because businesses are identical ex-ante, and because old and young consumers have the same, constant price elasticity, the markup is constant and equal to ν . The first-order condition for technology adoption equates its marginal cost with its marginal benefit. From the demand system described above, when the business takes \mathcal{M}_Y as given, the market share of young consumers has unit elasticity with respect to $a(j)$. Increasing a expands the market share of young consumers; each additional unit of sales generate profits equal to $1 - 1/\mu(j)$.

3.2 Equilibrium and comparative statics

An equilibrium of this economy is defined as follows.

Definition 1 (Equilibrium). *Let θ describe the demographic structure of consumers. Then given θ and all other structural parameters, an equilibrium is given by business-level markups $\{\mu(j)\}_{j=1}^J$, technology adoption choices $\{a(j)\}_{j=1}^J$ and aggregate markups \mathcal{M}_Y and \mathcal{M}_O that satisfy equations (2), (3), and (1).*

The equilibrium of this model exists, is unique, and is symmetric across businesses. To see why, note that because markups are the same across businesses (Equation 2), the equilibrium technology

adoption choices also are. Letting a denote this common technology adoption choice, the aggregate markups are given by:

$$\mathcal{M}_Y = a^{-(\nu-1)}\nu, \quad \mathcal{M}_O = \nu.$$

Therefore, the first-order condition characterizing a can be written as:

$$ac'(a) = \frac{\theta}{J} \left(1 - \frac{1}{\nu}\right).$$

The left-hand side is a strictly increasing bijection of $[1, +\infty)$ on $[0, +\infty)$, establishing the result.

Result 1 (Effects of demographics on technology diffusion). *The equilibrium technology adoption choice a is increasing in θ :*

$$\frac{da}{d\theta} > 0.$$

Thus the model implies that, when the share of young consumers is higher (that is, when θ increases), there is a stronger incentive for individual businesses to invest in technology, in order to gain market share.³⁵ In the following section, we will test this prediction by contrasting the technology adoption choices of firms that face consumer bases with different demographics.

The model also spells out how demographics can mediate the effects of changes in business' cost of adopting the technology. To simplify the discussion, we assume a particular functional form for the cost function $c(a)$:

$$c(a) = \frac{\gamma}{2}(a-1)^2, \quad \gamma > 0.$$

A change in the parameter γ could represent either an innovation that makes the technology easier or less costly to adopt for all businesses, or more generally, a change in the economic environment that affects the marginal cost of adoption for businesses (including, for instance, subsidies to adoption).

Result 2 (Demographics, adoption costs, and technology diffusion). *An increase in γ leads to a lower rate of technology adoption among businesses; and the effect is stronger, the higher the young share of the population, θ :*

$$\frac{da}{d\gamma} < 0, \quad \frac{d^2a}{d\gamma d\theta} < 0.$$

Clearly, following a reduction in the cost of adoption, the diffusion rate, as proxied by the equilibrium choice of a , increases. The main content of the result is that this effect is stronger when the customer base of businesses is younger (that is, when θ is higher). Conversely, if the customer base of businesses is older, then a change in the marginal cost of adoption need not necessarily lead to widespread diffusion of the technology.^{36,37}

³⁵Note that there are no countervailing effects on demand, because the markup is constant across demographic groups. If, instead, young consumers were also more price-elastic than old consumers, businesses might have a weaker incentive to adopt the technology, as this would increase their market share of the most price-elastic consumers. We do not include this heterogeneity in order to focus on the effects of differences in attitudes toward technology.

³⁶This result holds more generally, for any change in the cost function $c(a)$ that reduces the marginal cost of adoption for any level of adoption a beyond the specific cost function used here for illustrative purposes.

³⁷A point to note about the model's predictions is that total sales of businesses do not depend on demographics

In summary, starting from the assumption that young consumers’ taste are more sensitive to technology offerings, this model shows that technology adoption by businesses is stronger when their consumer base is younger. Additionally, the model suggests that a decline in the cost of adoption by businesses will generally have a weaker effect on overall adoption in contexts where the consumer base is older. We provide evidence of the former effect in the next section, before discussing the broader implications of the latter effect.

4 Evidence

This section uses business-level data on mobile payment adoption to test whether a merchant’s decision to adopt mobile payments is indeed influenced by the composition of its customer base.

4.1 Data

The data for our analysis comes from a prominent fintech company in India that caters to small and medium-sized businesses. This company provides businesses with physical terminals and digital payment management systems to facilitate the receipt and processing of payments across various networks. For our study, the dataset enables us to observe the decision of new stores to adopt one of the firm’s terminals and their subsequent usage patterns.³⁸ Our analysis will focus on examining how the adoption of the fintech company’s terminals by stores has evolved over time.³⁹

In particular, our study focuses on a shift in the types of payment services provided by the fintech company that occurred in 2019. Historically, the company had only offered traditional point-of-sale (POS) terminals, which required a physical card to conduct a transaction. Starting in May 2019, the company expanded its offerings to include mobile payment options through QR codes. This strategic shift was motivated by the increasing prevalence of mobile payments documented in Section 2. A merchant could still obtain a regular POS terminal after 2019: however, starting on May 2019, the fintech company started to also offer QR code enabled terminals, that would

(or on the cost of adoption). Because the equilibrium is symmetric, the market share of each business, whether of young or old consumers, is equal to $s_O(p(j)) = s_Y(p(j), a(j)) = 1/J$, and so total sales across businesses are simply equal to w . Thus in equilibrium, changes in demographics do not affect sales, even if they raise technology adoption by firms. This is for two reasons. First, because all businesses are identical, despite the fact that each of them attempts to attract more young customers by adopting the technology, their efforts cancel out. Second, there are no general equilibrium effects, because the wage w is fixed; in particular, more technology adoption does not increase wages and therefore demand for businesses’ products. Relaxing either assumption may lead to different predictions for movements in total sales.

³⁸In the data, a store is defined as a combination of one or more terminals owned by the same firm within a six-digit pincode. In other words, the assumption is that, if a firm owns multiple terminals in the same narrowly defined location, they are assumed to operate as part of the same store. To be clear, this assumption is unlikely to have any impact our analysis, because most firms in the data own only one terminal and operate in only one pincode.

³⁹We determine store adoption based on the date of first-time terminal usage provided by our fintech company. For a significant subset of the data, we also have information about the terminal installation date, enabling us to validate our primary adoption measure. Upon comparing our adoption time with the terminal installation month for the sample of terminals adopted in the sample period, we find that the two measures coincide exactly for almost 86% of the terminals (and this increases to over 94% when we allow for one period delay). This evidence validates our baseline approach.

allow individuals to directly use mobile payment options, for instance paying using UPI through any supporting apps.⁴⁰ Lastly, although our fintech company is sizable, it represents just one among various entities providing mobile payment solutions to merchants in India. Consequently, the decision to adopt QR-code payments is unlikely to significantly enhance consumer benefits from using UPI through network effects.⁴¹

Aside from the data provided by our fintech company, we also use public data on demographic and economic outcomes at the district level from the 2011 Census of India. Among other things, we use these data to construct measures of age structure for specific district, as well as other location-specific characteristics that allow us to adjust for other differences across areas in India (such as population, measures of economic activity, literacy, and others).⁴² Finally, we manually collected a list of universities in the country as of 2019, and mapped each university to its official pincode.⁴³ These data will be used in some of our validation analyses below.

4.2 Identification strategy

The model laid out in Section 3 shows that when consumers have different attitudes toward technologies, the distribution of these preferences should influence technology adoption by businesses. The prediction in the context of mobile payments is the following: merchants are likely to show greater interest in mobile payment technologies in areas with a higher concentration of young adults. In this section, we leverage data from our fintech payment company to test this prediction empirically.

To more accurately frame the empirical predictions of our model, we introduce the following ideal experiment. To start, we consider different groups of merchants and randomly allocate customer groups to each merchant group, with each customer group having a different age structure. This step aims to introduce exogenous variation in customer age, independent of merchants' characteristics. After this initial step, we would propose a dual offering where half of the merchant groups are presented with a traditional POS system exclusively for card transactions, while the other half are provided with terminals capable of mobile payments. We would then study how the adoption of the mobile-enabled terminal varies across groups as a function of the age of the consumer base. This experiment would allow us to estimate the extent to which variation in customer age could influence merchants' technology adoption decisions.

In our study, we emulate the experiment by using two sources of variation: technology avail-

⁴⁰In particular, the company offered both terminals that are enabled for both traditional cards and QR combined, as well as QR-code only terminals, that could be used only for mobile payments. Note that, in principle, a QR code could also be connected to a credit card. However, this option appears to be used very infrequently in our data, as most of QR transactions are UPI. Appendix A.2.2 contains more details on the data provided by our fintech partner and on the different POS offerings.

⁴¹Therefore, our context diverges from Agarwal et al. (2020)'s study on Singapore's largest bank introducing mobile payments and reducing cash usage. There, the involvement of the country's largest bank meant the shift prompted a significant change in the payment ecosystem.

⁴²A district is an administrative unit in India. There are 640 districts in the 2011 Census, with an average of 23 districts per state. There are about 2 million residents per district, which is close to the average population of a county in the United States.

⁴³Data Appendix A.2.3 discusses the construction of this data.

ability and client age demographics. The first source of variation comes from our company’s May 2019 launch of mobile payment options; this allows us to observe adoption rates at the same locations before and after the mobile payment option became available. The second source of variation comes from differing age demographics across Indian districts, enabling us to assess if an increase in adoption is related to the age of the potential customer base. Unlike the ideal experiment described above, however, age structure is not randomly assigned across Indian districts. Therefore, we will also need to convincingly show that our findings are attributable to age rather than other confounding factors that might influence adoption decisions.

The estimation of our empirical model would allow us to test the key prediction of the model: merchants facing more young consumers see a larger value in using mobile payments. This interpretation does not hinge on whether merchants that newly started a payment terminal from the fintech company were entirely new adopters of mobile payments, or transitioned from other mobile payment providers. In either case, observing higher adoption rates in younger districts would still be consistent with mobile payments being more valuable for these merchants.⁴⁴

4.3 Baseline specification

We implement this strategy by estimating a difference-in-differences model measuring how overall adoption of terminals of the fintech company increased after May 2019 across districts characterized by different age structures. The reduced-form model is:

$$y_{dt} = \alpha_d + \alpha_t + \beta (\text{AgeStructure}_d \times 1_{\{t \geq t_0\}}) + \Gamma'_t \mathbf{X}_d + \epsilon_{dt}, \quad (4)$$

where y_{dt} is a measure of adoption of the firm’s terminals in district d and in month t ; α_t and α_d are respectively time and district fixed-effects. AgeStructure_d in our baseline model is the share of adults (i.e., 15-74 years old) that are less than 30 years of age, according to the 2011 Census — the most recent Census before the technology rollout, but we consider also alternative treatment definitions below. We always z -score the treatment variable to facilitate the comparison across variables. \mathbf{X}_d is a vector of district characteristics, measured before the policy, and is allowed to have time-varying effects on the outcome, Γ_t . In all our analyses, we use six months of data before May 2019 (i.e., pre-period) and six months after.⁴⁵ When plotting the dynamic effects, we normalize the last month of the pre-period (i.e., April 2019) to zero in the following specification:

$$y_{dt} = \alpha_d + \alpha_t + \sum_{k=-6, k \neq -1}^{k=+6} \beta_k (\text{AgeStructure}_d \times 1_{\{t=t_0+k\}}) + \Gamma'_t \mathbf{X}_d + \epsilon_{dt}. \quad (5)$$

⁴⁴If anything, the availability of other mobile-enabled payment systems in the district, and the awareness with merchants regarding these systems, should bias us toward documenting weaker adoption responses.

⁴⁵To be clear, May 2019 is considered as a treated month, as the announcement and formal initiation of the policy occurred in May. In dynamic specifications, we directly model the effect for this month and find that the impact in May is generally null. This aligns with the observation that this month received only partial treatment and furthermore reflects the company’s initial minimal effort to acquire customers to ensure smooth integration of the new mobile option into the ecosystem.

In our baseline specification, we measure the level of adoption y_{dt} with the number of new stores that obtained a terminal from the firm in that month, scaled by the number (in hundreds) of firms in the district from the Census.⁴⁶ However, as robustness, we also consider alternative ways to measure the same outcome, which we discuss below. Standard errors are clustered at the district level (Bertrand et al. 2004).

4.4 Main Results

Table 2 presents the results from the baseline specification 4. On average, our firms experienced a larger increase in new businesses joining the platform in areas with a younger population, consistent with our initial hypothesis (column 1). Specifically, we find that one standard deviation increase in the share of young population led to almost 0.05 new businesses joining the platform per hundred firms in the district. This corresponds to roughly a 25% increase relative to the adoption rate right before the policy change (i.e., April 2019).

Panel (a) of Figure 3 reproduces the same finding using the dynamic specification, which allows us to identify changes in adoption month-by-month. Consistent with the validity of our design, we find that the share of young adults is not connected in a significant way with adoption during the pre-period. Furthermore, after May 2019, we see a significant increase in adoption. The effect increases over time, with the effect size peaking at more than 0.1 new stores per hundred firms in the district. Panel (a) of Appendix Figure A-5 reproduces the same analysis using the inverse hyperbolic sine transformation of the number of new stores joining the platform as the outcome variable. The overall pattern is similar: districts with more young adults do not outperform in their adoption of the new technology before May 2019, but saw an adoption spike after.⁴⁷

A key concern with our analysis is that age differences are likely to correlate with other district characteristics and these factors may also potentially influence the impact of the policy change on adoption. In principle, this consideration is not inconsistent with our overall approach: if younger individuals are more inclined towards mobile payment options, this should coincide with other distinctions in the local economy. Nonetheless, it is crucial to ensure that our findings primarily reflect the impact of age demographics rather than ancillary factors.

This issue is evident in Table 3, where we observe that districts with a higher proportion of young adults differ significantly across various dimensions. For example, these areas are generally smaller, exhibit lower literacy rates, have fewer schools, have a reduced percentage of the working

⁴⁶To clarify, our analysis covers the overall adoption of our firm’s products, not just QR-enabled ones. This method is justified for several reasons. Firstly, our empirical approach focused solely on QR-enabled terminals would be impossible, as their presence was nonexistent before the period in question (though we will later conduct a post-period test on this variable, Appendix Table A-1). Secondly, as explained in the thought experiment, the examination of how overall adoption evolves over time provides insights into how the addition of a mobile payment option changed the demand for merchants. In fact, our approach uses the info on adoption before May 2019 as a benchmark for the district demand for our Company’s product. Lastly, it is also useful to note that (as expected) the majority of adoption growth in the post-period is attributed to QR-enabled terminals: in fact, about 80% of the increase in new platform members is due to stores opting for QR-enabled terminals.

⁴⁷The same result is also presented in Table 2, column 3.

population, and are less densely populated, among other traits. Notably, districts with a younger population also tend to feature fewer stores using our partner company’s services and record fewer transactions on their platform.⁴⁸

Before addressing this issue empirically, it is critical to highlight that the direction of most relationships observed is actually the opposite of what one might anticipate if an omitted variable were explaining the positive link between age structure and adoption. The prevailing literature on technology adoption usually indicates that newer technologies are more readily adopted in areas with higher education levels and greater wealth (Caselli and Coleman 2001). Contrary to this, our findings imply that, if anything, regions with a younger population tend to be less educated and exhibit lower economic activity. This relationship is likely explained by the well-documented negative correlation between economic prosperity and fertility rates (Jones et al. 2010). If this observation is correct, we should find that including these controls increases the size of the effect.

To directly mitigate this concern, we incorporate a wide set of district-level controls in our empirical model. This approach helps us to net out the effects of other district characteristics from our desired treatment effects. In particular, our analysis controls for population, number of firms, the share of agricultural workers, the share of literate individuals, the share of the working population, and the average amount of night light in 2018, as a proxy of overall economic activity.⁴⁹ As we show in Table 3, we find that once we control for these variables, districts with different demographic structures do not differ across other observed characteristics.

Panel (b) of Figure 3 reports our main figure including the controls interacted with month dummies. Consistent with our intuition, the magnitude of the effects increases slightly when controls are included. This difference in magnitude can be better appreciated when estimating a single parameter, capturing the overall average effect size (Table 2, columns 2 and 4). However, the general conclusion from the test is unchanged. We note that the same results are confirmed in the analysis estimating the effect on the percentage change in adoption (Panel (b) of Appendix Figure A-5). Importantly, our results are not driven by any single controls, therefore assuaging any concern on the robustness of the control selection. In Appendix Figure A-6, we plot the dynamic effects including one control at the time. Although the precise magnitude varies slightly across specifications, both the sign and the general scale of the estimates remain stable across models.

Last, we show that the age structure not only predicts the increase in adoption for our fintech company following the introduction of QR-enabled terminals but also explains the share of stores that opted for QR-enabled terminals in the post-period. Indeed, stores could continue to adopt our fintech company’s services and request card-only POS terminals after May.⁵⁰ In Table A-1,

⁴⁸In principle, this stylized fact is consistent with our theory: before May 2019, our partner company did not offer mobile options and therefore this company was less attractive in areas where a larger share of the population has a preference for mobile payments.

⁴⁹Variables that are aggregates (i.e., population) are included after being log-transformed. We selected controls in a parsimonious way, and we discuss below (Appendix Figure A-6) how our result is robust to alternative ways to select the set of controls included in this analysis.

⁵⁰As we mentioned before, while some firms continue adopting card-only POS, we also see in aggregate that about 80% of the increase in new platform members is due to stores opting for QR-enabled terminals.

we explore whether the district-level share of adopting stores with QR-enabled terminals in the post-period is associated with the age structure. It is important to note that this analysis is strictly cross-sectional, as QR-enabled terminals were only available in the post-period. We find that districts with a younger population showed a higher preference for QR-enabled terminals among adopters, and this result holds both with and without our standard set of controls. This result aligns well with our model’s prediction, supporting the notion that the increase in adoption documented earlier was driven by the introduction of QR-enabled terminals.

Before concluding, we present some ancillary results. To start, we show that our results are robust to the treatment definition. In particular, Appendix Figure A-7 reproduces our main analysis with controls using the share of the total population less than 30 years old as treatment. In other words, our age structure index now includes also very young individuals (i.e., less than 15 years old), which were excluded from the baseline to allow for the possibility that this group is less likely to capture potential shopping customers and use electronic payments. Appendix Figure A-8 repeats the same exercise but defines the share of young adults focusing on those below 40 years old. In both cases, we standardize the treatment variable to have a mean of zero and a standard deviation of one, thus facilitating comparison across figures. In general, the results we obtain are almost identical: if anything, the magnitude of the effects is slightly larger. Furthermore, in Appendix Figure A-9, we show our main results when the outcome is the number of new stores joining the platform scaled by the population size (in 100,000s) rather than the number of firms. The scale of the coefficient is different, but the message remains unchanged.

Lastly, we examine whether the increase in adoptions leads to an overall surge in the number of stores active on the platform. In other words, rather than looking at new adoptions in a month, our outcome is now the total number of stores that have a terminal with our company, irrespective of whether they joined that month or earlier. Appendix Figure A-10 presents the result: indeed, the relatively larger increase in adoption in younger districts also translated into more stores active in the platform. The effect is sizable, as a one-standard-deviation increase in the treatment variable leads to about two extra stores per hundred firms in the district. This evidence confirms that the effect of adoption led to an overall increase in the business managed by the fintech company.

4.5 Age and historical fertility: a 2SLS application

As previously mentioned, one possible concern with our findings is that omitted variables influence both the age distribution and adoption rates. Specifically, younger individuals might gravitate towards areas of greater economic dynamism, which could independently drive the uptake of new digital payment methods, regardless of demographic disparities. Although this hypothesis seems at odds with the summary statistics presented earlier (Table 3), we cannot categorically dismiss the possibility that such a mechanism might operate through factors not observable in our data.

To address the concern regarding the selection of younger individuals into more economically dynamic areas, we introduce a test designed to isolate variation in the age distribution at the time of our study that is independent of migration patterns from previous decades and only captures

the historical characteristics of the districts. The underlying idea is as follows: the proportion of young adults in a given region is influenced by both migration trends and the fertility rates within districts several decades earlier. All else equal, it is reasonable to expect that districts that exhibited higher fertility rates in the 1990s will have a greater proportion of young adults by 2010. Hence, by leveraging only the variation in the current age structure attributable to differences in historical fertility rates, our findings should be shielded from critiques pertaining to migration effects.

The sex ratio from the past represents a good candidate for the current age structure. A skewed sex ratio can affect the marriage market and consequently fertility (Guilmoto 2012; Dyson 2012; Angrist 2000). Therefore, we should expect that regions with a more skewed sex ratio in the early 1990s may end up with fewer kids, and therefore a smaller share of young adults in the early 2010s. This idea seems to be supported by the data: in Appendix Figure A-11, we plot the share of young adults (i.e., adults less than thirty) in 2010 against the district-level sex ratio in 1991, measured as the ratio of male to female.⁵¹ Consistent with this idea, we find that districts that are at either tail of the distribution tend to have a smaller share of young adults twenty years later. This relationship can be confirmed formally: in Table 4, we predict the share of young adults used in this paper with the quadratic function of the sex ratio in 1991. The analysis finds that the historical sex ratio strongly predicts the future share of young adults, with the largest share of the young population present in districts with a sex ratio slightly above one.⁵² The Sanderson and Windmeijer (2016) multivariate first-stage F -statistics for the validity of the instruments is 43.46. As the Stock–Yogo 10 percent and 15 percent critical values for a perfectly identified model with two excluded instruments are, respectively, 19.93 and 11.59, we can reject that the instruments are weak.

Building on this result, we implement a 2SLS estimator, where we instrument the share of young adults in 2011 – our main treatment variable in the analyses above – with a quadratic function of the sex ratio in 1991. Before showing the result, we want to clarify what the purpose of this 2SLS is. Our goal is to replicate our main findings with a historical measure that is less likely to be affected by the level of dynamism in the district in 2019. We implement this approach as a 2SLS (rather than in reduced form) because this allows us to generate estimates that are directly comparable to the OLS presented before, while using historical information about the district. However, we recognize that an exclusion restriction is unlikely to hold, since the historical sex ratio may affect other aspects of the local economy beyond the age distribution.

With these caveats in mind, the results are presented in Figure 4: as usual, we present the results dynamically around May 2019. The estimates using the 2SLS are qualitatively close to our baseline estimates in terms of dynamics and magnitude (panel a): we find that the age structure does not predict differential adoption before May 2019, and we confirm that districts with more

⁵¹We source this information from the 2011 Census which provides the number of males and females across Indian districts by each decade since 1901. Importantly, the Census data provides this information using the definition of districts in 2011.

⁵²Our analysis suggests that the sex ratio maximizing the share of young adults in our context is between 1.1 and 1.2, with lower natality at the tail. This evidence appears consistent with the previous literature on the topic (Hesketh and Xing 2006; Hesketh and Min 2012).

young individuals saw a larger increase in adoption afterwards. However, the magnitude of the effect is larger. The same result also holds when we include all controls that were employed in our main specification, as considered before (panel b). Columns 2 and 3 of Table 4 confirms the same result when comparing the average behavior pre- and post-May 2019.

This evidence confirms that our results do not simply reflect the sorting of young people towards more dynamic areas. Instead, this evidence is consistent with the idea that structural demographic characteristics may be important to understand the diffusion of technologies from the business side. In fact, locations that were expected to have a higher share of young people based on historical demographic structure saw larger adoption after the QR code was available.

4.6 An alternative approach: the location of universities

The findings detailed above validate the key prediction of the model: merchants in regions populated by younger individuals tend to be more interested in adopting the new technology, mobile payment. As previously outlined, our interpretation of these findings is that a higher concentration of young adults amplifies the pool of potential customers with a preference for mobile over traditional card payments, thus boosting merchants' motivation to adopt mobile payment solutions. As a way to further bolster our result, we present a new test that does not rely on district-level measures of demographic structure, but exploits variation of consumer demand *within* a district.

In particular, we leverage the presence of universities in the country as a way to create differences in demand from young adults across neighborhoods within the same district. The premise is that neighborhoods with a university tend to receive a daily influx of young adults, who could constitute a significant share of customers. Concurrently, these areas are experiencing similar general economic and social conditions to those of businesses located in the same district but in a different neighborhood.⁵³ If this assumption holds, then a comparative analysis of adoption rates within a district, between areas with and without a university, could provide supplementary evidence to the discussion above and help address the issue raised.

Among the others, this test can help addressing specific concerns about the importance of the heterogeneity in business owners' characteristics in explaining our findings. In fact, a possible alternative interpretation is that regions with a younger demographic could also have a higher share of younger entrepreneurs who may be more inclined to adopt mobile payment methods, regardless of customer demand. Our university test can take care of this issue because neighborhoods with a university should face stronger demand from young adults, because of the daily influx of students. However, there is no basis to believe that business owners in these areas are systematically younger than those living in other areas of the same districts. Furthermore, this setting allows us to directly rule out this alternative interpretation by exploiting merchant-level variation. For instance, as we discuss below, we can use the set of businesses that do not normally deal with university students as a placebo test.

⁵³For instance, there is no basis to believe that business owners in university areas are younger or belong to a different social group than those in other parts of the district.

To test this hypothesis, we manually compiled a list of universities in India and linked each to its official location, identified by a (six-digit) pincode.⁵⁴ With this data, we are able to identify the list of pincodes in India that host at least one university. Then, we estimate a differences-in-differences estimator of the following form:

$$y_{pt} = \alpha_{dt} + \alpha_p + \sum_{k=-6, k \neq -1}^{k=+6} \gamma_k (1\{Univ = 1\}_p \times 1_{\{t=t_0+k\}}) + \nu_{pt} \quad (6)$$

where y_{pt} is a measure of adoption of terminals provided by our fintech company at the pincode p and month t level; α_{dt} and α_p are, respectively, district-by-month and pincode fixed-effects; $1\{Univ = 1\}_p$ is a dummy variable equal to one if the pincode has at least one university. Similar to before, we use six months of data before May 2019 (i.e., pre-period) and six months after. Therefore, the last month of data used in this analysis is November 2019. When plotting the dynamic effects, we normalize the last month of the pre-period (i.e., April 2019) to zero. Given the absence of information on the number of firms that operate in a pincode, we cannot utilize the rate of adoption relative to firms as our primary outcome y_{pt} . Nonetheless, in line with the findings presented earlier, we will employ both the raw number of adoptions and the inverse hyperbolic sine (IHS) transformation of the adoption numbers as alternative measures. Standard errors are clustered at the pincode level to accommodate this approach.

The findings are detailed in Figure 5: in line with our initial hypothesis, we observe that pincodes hosting a university experienced a more substantial increase in the number of stores adopting our fintech services compared to other pincodes within the same district. The magnitude of the effects is significant: pincodes with a university witnessed approximately a 20% greater increase in adoption over a few months, with this disparity enduring throughout most of the post-intervention period. Appendix Figure A-12 corroborates this result by utilizing the raw count of new stores adopting the fintech company’s terminals each month as the outcome.⁵⁵ Table 5 confirms the same result when comparing the average behavior pre- and post-May 2019.

Therefore, the presence of university students plays a significant economic role in explaining the increased demand for mobile payment options in our sample. We interpret these findings as evidence that young adult customers—who tend to prefer mobile payments over cards, as documented earlier—affect local businesses’ adoption of mobile payment technology. To reinforce this interpretation, we now conduct a set of tests exploiting variation in merchant types. One advantage of the university analysis is that it assumes the increase in demand originates specifically from university students. This feature allows us to generate predictions about the heterogeneity of the

⁵⁴Data Appendix A.2 outlines the methodology employed to gather this information. It’s important to acknowledge that the pincode data primarily reflects the headquarters or the main building of the university. For larger institutions, certain facilities might reside beyond this designated location. Nonetheless, as elaborated in the Appendix, we anticipate that this detail will not pose a substantial issue. If it does have any impact, we expect that it would bias our findings towards null effects in our analysis.

⁵⁵One may be concerned that most districts do not have any university, and therefore, generate no useful variation in the analysis. As a sanity check, in Appendix Figure A-13 we replicate the findings discussed above by manually dropping districts without any university and find identical results.

effects at store level. In particular, we expect an increase in mobile payment adoption, specifically among merchants who are likely to interact with students.

To identify businesses that typically serve students, we first look at those businesses in the non-tradable sector that would generally depend on very localized demand from consumers. In our data, these are made up of retailers, gas stations, restaurants, leisure facilities, personal services, and transportation. Additionally, we also consider an alternative approach which broadens this category to also include financial services, healthcare (e.g., pharmacies), and educational services.⁵⁶ We then replicate our main results using only stores that belong to these categories. The findings, presented in columns 1 and 2 of the two panels of Table A-2, support our hypothesis: the increase in mobile payment adoption in university areas is predominantly from businesses within these sectors that are likely serving students on a regular basis. Indeed, the results from these sub-samples are positive and qualitatively similar to our main findings. Moreover, as before, the surge in adoption is solely attributed to changes post-May 2019 (panels (a) and (b) in Figure A-14).

To validate this idea, we also executed a complementary analysis on a set of businesses likely indifferent to local consumer demands in their decision to adopt mobile payment methods. These include government and regulated sectors, manufacturing, wholesale, warehouse operations, and professional services. This test serves as a placebo; if our model accurately captures how business technology decisions reflect student demand, we should observe no significant effects in these sectors. Conversely, finding an impact here could indicate that our results are influenced by other business owner characteristics that vary by university presence. The results confirmed our hypothesis: university presence did not affect mobile adoption in these sectors, with the effects being not only statistically insignificant but also small in magnitude and precisely estimated (panel (c) in Figure A-14).⁵⁷

This evidence collectively confirms that the presence of young adults – represented here by students — is a crucial determinant of merchants’ decisions to adopt point-of-sale systems compatible with mobile payments.⁵⁸

⁵⁶To be precise, we define consumer facing merchants as merchants categorized as clothing and accessories, consumer durable, retail consumer goods, restaurants and hotels, gas station, personal services (e.g., hair salons), telecom services, transportation. We incorporate professional services by adding businesses that are categorized as health, financial, and education.

⁵⁷In columns 4 of the panels of Table A-2, we also look at those merchants that could not be categorized in any of the groups discussed. In this sub-sample, we find a positive effect, consistent with the main findings. This result is not surprising for us, because the sample of stores that are not categorized is mostly made-up businesses that belong to a “miscellaneous” category in our data (91%). We expect this category to be mostly made up by very small establishment that belong broadly to the retail sector, while failing to fit clearly in one of its sub-categories (e.g., street cart serving prepared food but also selling produces).

⁵⁸Our preferred interpretation is that these results stem from the young age of the students, consistent with other demographic results shown earlier. However, we also recognize that students may differ systematically in other dimensions besides age, making this interpretation not the only possible explanation. Despite this limitation, the tests still confirm the broader point: differences in consumer composition can drive varying adoption rates across businesses.

5 Conclusion

In this paper, we asked whether demographics can influence the rate at which new technologies diffuse in an economy. We studied the particular case of mobile payment technology in India. We started by noting that empirically, age is a central determinant of the propensity to use mobile payments in India. We then showed that, in a simple model of technology adoption, this fact implies that businesses are more likely to adopt the technology when facing a young customer base. Finally, we used data on the adoption of mobile-enabled payment terminals by a leading Indian fintech to show that Indian locations in which the population is younger exhibit a higher propensity to adopt the technology once it is introduced.

Thus the core message of our analysis is that younger consumers exhibit markedly different preferences for mobile payments in India, and that these customer preferences shaped merchants' decision to adopt the mobile payment technology we analyzed in this paper. While our paper focuses specifically on mobile payments in India, we believe that some of the implications of our analysis may be broader. In fact, the idea that younger individuals may be more responsive to new technologies – the underlying assumption in our model – is likely to hold across different countries and technologies. For instance, Figure A-15 examines the use of Pix in Brazil. Pix is a form of payment that shares many similarities to UPI (Sarkisyan 2023). Consistent with our evidence, we find that areas with a younger population shows stronger penetration of this technology.⁵⁹

In this context, we believe that there are two broad conclusions from this paper. First, in the context of mobile payments, an interesting question is whether the dramatic speed of diffusion in India is tied to the particular demographic structure of the country. India has a significantly younger population than developed nations, with a median age of 28, compared to 40 in OECD countries.⁶⁰ Our analysis suggests that this peculiarity may indeed have played a role.

Second, our results speak to the broader question of what determines the rate of diffusion of new technologies. We highlight that for consumer-facing technologies, demographics appear to play a large quantitative role, both directly, and indirectly, through their impact on business incentives. This result raises the possibility that population aging leads to slower rates of technology diffusion. The connection between age and diffusion has potential implications for understanding how policy can, or cannot, spur the adoption of new technologies. Additionally, for technologies where adoption decisions are complements (such as network-based technologies), age effects may interact with, and potentially amplify, coordination problems in adoption. We leave these issues for future research.

⁵⁹We thank Sergey Sarkisyan to make this point, providing the first version of this analysis when discussing our paper. We partially modified the initial version he created to make this test closer to our own approach.

⁶⁰<https://www.oecd-ilibrary.org/sites/d56a2fbc-en/index.html?itemId=/content/component/d56a2fbc-en>

References

- Abeliansky, A. L. and K. Prettner (2023). Automation and population growth: Theory and cross-country evidence. *Journal of Economic Behavior & Organization* 208, 345–358.
- Acemoglu, D. and P. Restrepo (2017). Secular stagnation? the effect of aging on economic growth in the age of automation. *American Economic Review* 107(5), 174–179.
- Acemoglu, D. and P. Restrepo (2022). Demographics and automation. *The Review of Economic Studies* 89(1), 1–44.
- Agarwal, S., S. Alok, P. Ghosh, and X. Zhang (2022). How does access to credit impact consumer expenditure and savings?
- Agarwal, S., P. Ghosh, J. Li, and T. Ruan (2024). Digital payments and consumption: evidence from the 2016 demonetization in india. *The Review of Financial Studies*.
- Agarwal, S., W. Qian, Y. Ren, H.-T. Tsai, and B. Y. Yeung (2020). The real impact of fintech: Evidence from mobile payment technology. *Available at SSRN 3556340*.
- Agarwal, S., W. Qian, B. Y. Yeung, and X. Zou (2019). Mobile wallet and entrepreneurial growth. In *AEA Papers and Proceedings*, Volume 109, pp. 48–53.
- Aggarwal, B., N. Kulkarni, and S. Ritadhi (2023). Cash is king: The role of financial infrastructure in digital adoption. *The Review of Corporate Finance Studies* 12(4), 867–905.
- Aksoy, Y., H. S. Basso, R. P. Smith, and T. Grasl (2019). Demographic structure and macroeconomic trends. *American Economic Journal: Macroeconomics* 11(1), 193–222.
- Alok, S., P. Ghosh, N. Kulkarni, and M. Puri (2024). Open banking and digital payments: Implications for credit access.
- Anelli, M., G. Basso, G. Ippedico, and G. Peri (2019). Youth drain, entrepreneurship and innovation. Technical report, National Bureau of Economic Research.
- Angrist, J. D. (2000). Consequences of imbalanced sex ratios: Evidence from america’s second generation. *NBER Working Paper* (w8042).
- Atkin, D., A. Chaudhry, S. Chaudry, A. K. Khandelwal, and E. Verhoogen (2017). Organizational barriers to technology adoption: Evidence from soccer-ball producers in pakistan. *The Quarterly Journal of Economics* 132(3), 1101–1164.
- Azoulay, P., B. F. Jones, J. D. Kim, and J. Miranda (2020, March). Age and high-growth entrepreneurship. *American Economic Review: Insights* 2(1), 65–82.
- Bachas, P., P. Gertler, S. Higgins, and E. Seira (2021). How debit cards enable the poor to save more. Technical report.
- Bertrand, M., E. Dufflo, and S. Mullainathan (2004). How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics* 119(1), 249–275.
- Bian, W., L. W. Cong, and Y. Ji (2023). The rise of e-wallets and buy-now-pay-later: Payment competition, credit expansion, and consumer behavior. Technical report, National Bureau of Economic Research.

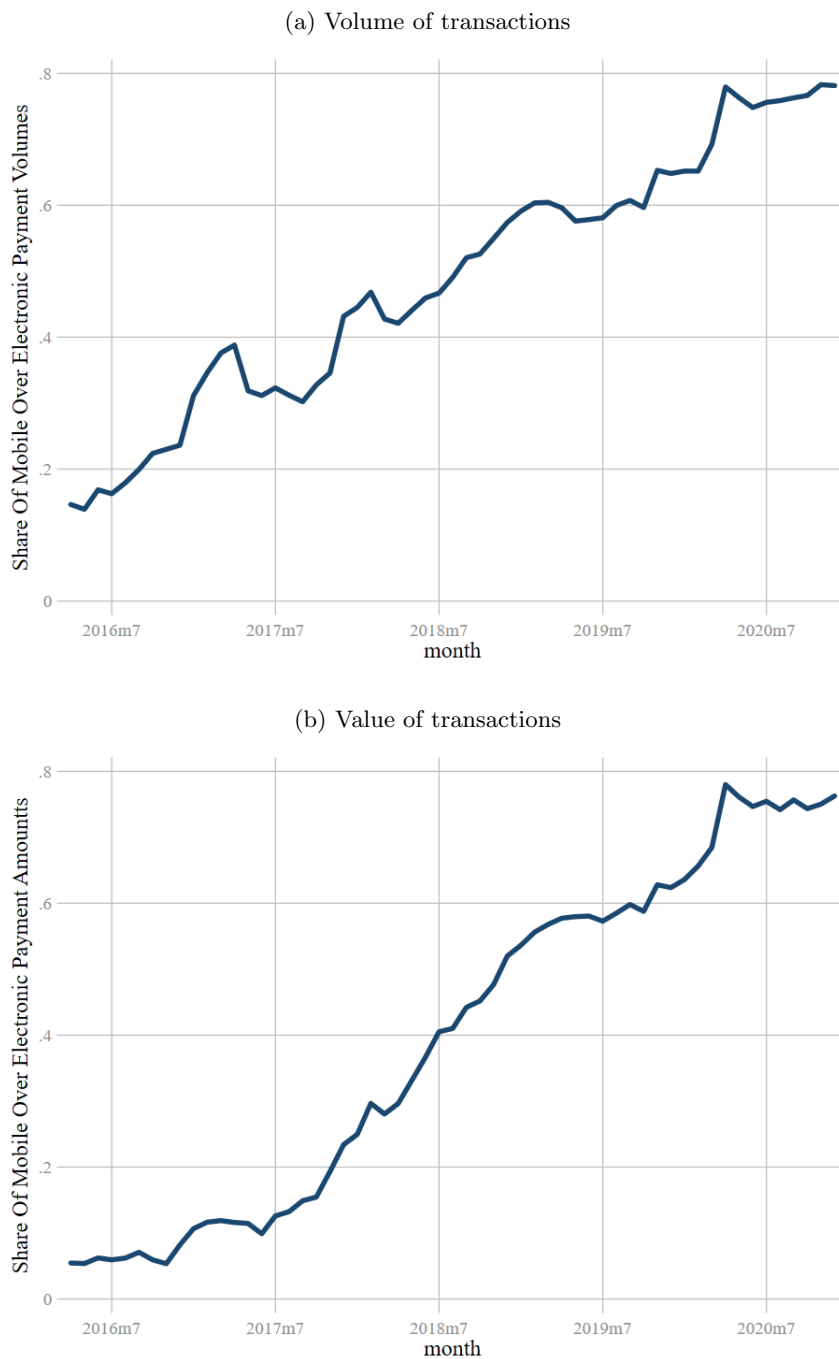
- Biasi, B. and S. Ma (2022). The education-innovation gap. Technical report, National Bureau of Economic Research.
- Bloom, D., D. Canning, and J. Sevilla (2003). *The demographic dividend: A new perspective on the economic consequences of population change*. Rand Corporation.
- Bornstein, G. (2021). Entry and profits in an aging economy: The role of consumer inertia. Technical report, working paper.
- Caselli, F. and W. J. Coleman (2001). Cross-country technology diffusion: The case of computers. *American Economic Review* 91(2), 328–335.
- Cheng, W. and B. A. Weinberg (2024). Old and connected versus young and creative: Networks and the diffusion of new scientific ideas. Technical report, National Bureau of Economic Research.
- Chodorow-Reich, G., G. Gopinath, P. Mishra, and A. Narayanan (2019). Cash and the Economy: Evidence from India’s Demonetization. *The Quarterly Journal of Economics* 135(1), 57–103.
- Comin, D. and B. Hobijn (2010). An exploration of technology diffusion. *American economic review* 100(5), 2031–2059.
- Conley, T. G. and C. R. Udry (2010). Learning about a new technology: Pineapple in Ghana. *American Economic Review* 100(1), 35–69.
- Costinot, A., D. Donaldson, M. Kyle, and H. Williams (2019). The more we die, the more we sell? a simple test of the home-market effect. *The quarterly journal of economics* 134(2), 843–894.
- Crouzet, N., A. Gupta, and F. Mezzanotti (2023). Shocks and technology adoption: Evidence from electronic payment systems. *Journal of Political Economy* 131(11), 3003–3065.
- Das, S., L. Gadenne, T. K. Nandi, and R. Warwick (2022). Does going cashless make you tax-rich? evidence from India’s demonetization experiment.
- Derrien, F., A. Kecskés, and P.-A. Nguyen (2023). Labor force demographics and corporate innovation. *The Review of Financial Studies* 36(7), 2797–2838.
- Dubey, T. S. and A. Purnanandam (2023). Can cashless payments spur economic growth? Available at SSRN 4373602.
- Dynan, K. E. (2000). Habit formation in consumer preferences: Evidence from panel data. *American Economic Review* 90(3), 391–406.
- Dyson, T. (2012). Causes and consequences of skewed sex ratios. *Annual Review of Sociology* 38, 443–461.
- Feyrer, J. (2007). Demographics and productivity. *The Review of Economics and Statistics* 89(1), 100–109.
- Feyrer, J. (2008). Aggregate evidence on the link between age structure and productivity. *Population and Development Review* 34, 78–99.
- Foster, A. D. and M. R. Rosenzweig (2010). Microeconomics of technology adoption. *Annual Review of Economics* 2(1), 395–424.
- Goehring, G., F. Mezzanotti, and S. A. Ravid (2023). Technology adoption and career concerns: Evidence from the adoption of digital technology in motion pictures.

- Gordon, R. J. (2017). *The Rise and Fall of American Growth*. Princeton: Princeton University Press.
- Guilmoto, C. Z. (2012). Skewed sex ratios at birth and future marriage squeeze in china and india, 2005–2100. *Demography* 49(1), 77–100.
- Gupta, A., J. Ponticelli, and A. Tesei (2022). Language barriers, technology adoption and productivity: Evidence from agriculture in india. *Working Paper*.
- Hall, B. H. and B. Khan (2003). Adoption of new technology. Technical report, National bureau of economic research.
- Hashimoto, K.-i. and K. Tabata (2016). Demographic change, human capital accumulation and r&d-based growth. *Canadian Journal of Economics/Revue canadienne d'économique* 49(2), 707–737.
- Hesketh, T. and J. M. Min (2012). The effects of artificial gender imbalance: Science & society series on sex and science. *EMBO reports* 13(6), 487–492.
- Hesketh, T. and Z. W. Xing (2006). Abnormal sex ratios in human populations: causes and consequences. *Proceedings of the National Academy of Sciences* 103(36), 13271–13275.
- Higgins, S. (2023). Financial technology adoption: Network externalities of cashless payments in mexico.
- Hu, A. S., C. A. Parlour, and U. Rajan (2019). Cryptocurrencies: Stylized facts on a new investible instrument. *Financial Management* 48(4), 1049–1068.
- Huettner, F. and M. Sunder (2012). Rego: Stata module for decomposing goodness of fit according to owen and shapley values. In *United Kingdom Stata Users' Group Meetings 2012*, Number 17. Stata Users Group.
- Israeli, O. (2007). A shapley-based decomposition of the r-square of a linear regression. *The Journal of Economic Inequality* 5, 199–212.
- Jack, W. and T. Suri (2014). Risk sharing and transactions costs: Evidence from kenya's mobile money revolution. *American Economic Review* 104(1), 183–223.
- Jiang, E. X., G. Y. Yu, and J. Zhang (2024). Bank competition amid digital disruption: Implications for financial inclusion.
- Jones, L. E., A. Schoonbroodt, and M. Tertilt (2010). Fertility theories: Can they explain the negative fertility-income relationship? In *Demography and the Economy*, pp. 43–100. University of Chicago Press.
- Klee, E. (2008). How people pay: Evidence from grocery store data. *Journal of Monetary Economics* 55(3), 526–541.
- Lewis, E. (2011). Immigration, skill mix, and capital skill complementarity. *The Quarterly Journal of Economics* 126(2), 1029–1069.
- Liang, J., H. Wang, and E. P. Lazear (2018). Demographics and entrepreneurship. *Journal of Political Economy* 126(S1), S140–S196.

- Ludwig, A., T. Schelkle, and E. Vogel (2012). Demographic change, human capital and welfare. *Review of Economic Dynamics* 15(1), 94–107.
- Maestas, N., K. J. Mullen, and D. Powell (2023). The effect of population aging on economic growth, the labor force, and productivity. *American Economic Journal: Macroeconomics* 15(2), 306–332.
- Manuelli, R. E. and A. Seshadri (2014). Frictionless technology diffusion: The case of tractors. *American Economic Review* 104(4), 1368–1391.
- Mezzanotti, F. and T. Simcoe (2023). Innovation and appropriability: Revisiting the role of intellectual property. Technical report, National Bureau of Economic Research.
- Munshi, K. (2004). Social learning in a heterogeneous population: technology diffusion in the indian green revolution. *Journal of Development Economics* 73(1), 185–213.
- Peters, M. and C. Walsh (2019). Declining dynamism, increasing markups and missing growth: The role of the labor force. Technical report, Society for Economic Dynamics.
- Prensky, M. (2001a). Digital natives, digital immigrants part 1: Do they really think differently? *On the horizon*.
- Prensky, M. (2001b). Digital natives, digital immigrants part 2: Do they really think differently? *On the horizon*.
- Sanderson, E. and F. Windmeijer (2016). A weak instrument f-test in linear iv models with multiple endogenous variables. *Journal of Econometrics* 190(2), 212–221.
- Sarkisyan, S. (2023). Instant payment systems and competition for deposits.
- Suri, T. and W. Jack (2016). The long-run poverty and gender impacts of mobile money. *Science* 354(6317), 1288–1292.
- Vallee, B., P. Ghosh, and Z. Yao (2024). FinTech Lending and Cashless Payments. *The Journal of Finance*.
- Wang, Z. and A. L. Wolman (2016). Payment choice and currency use: Insights from two billion retail transactions. *Journal of Monetary Economics* 84, 94–115.
- Yermack, D. (2018). Fintech in sub-saharan africa: What has worked well, and what hasn't. Technical report, National Bureau of Economic Research.


Figures and Tables

Figure 1: Share of electronic payments done using a mobile option



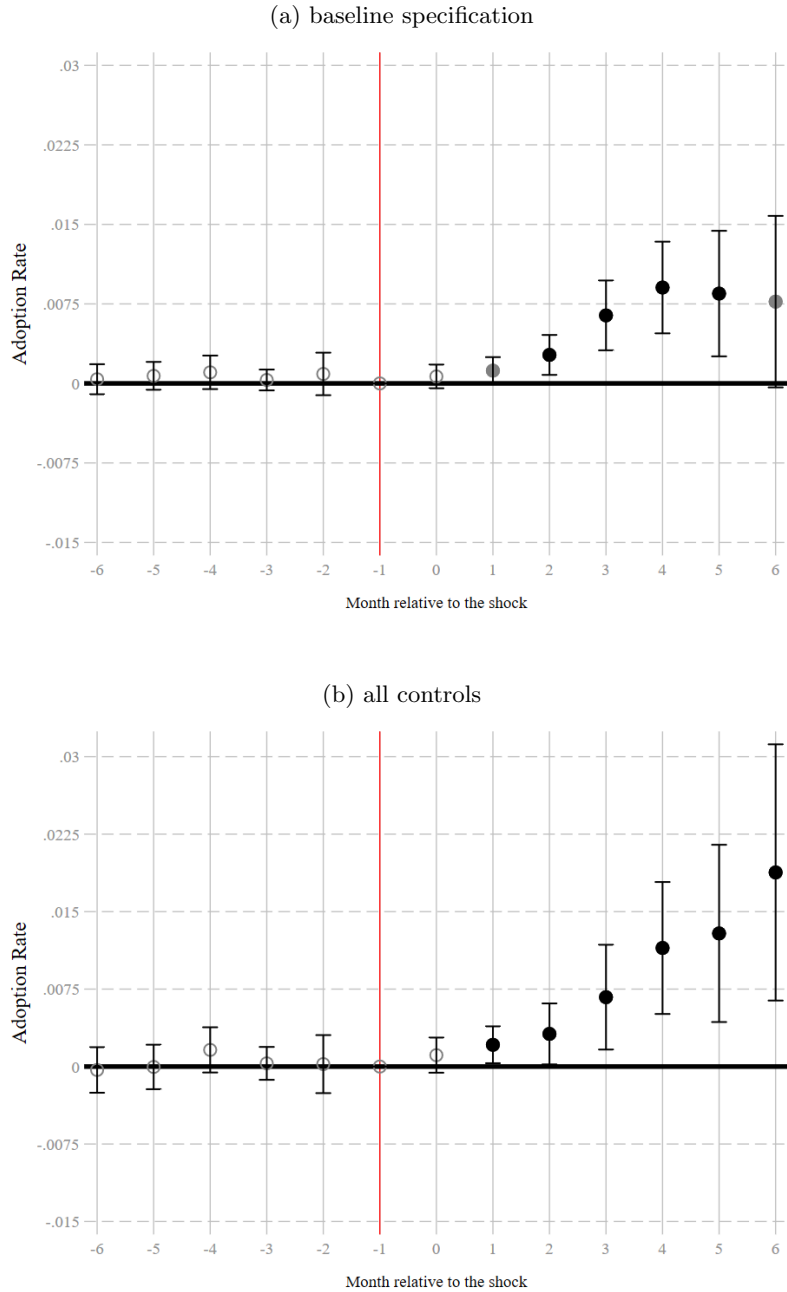
Notes: These two figures plot the share of electronic transactions that are done using mobile (at monthly level). Panel (a) reports the measure based on volume of transactions, while panel (b) examines the value of transaction. Electronic transactions are defined as the sum of UPI, mobile wallets, and cards (debit, credit, and prepaid), excluding the use of cards at the ATM. Mobile transactions are defined as the sum of UPI and mobile wallets. The data to construct these figures come from the Reserve Bank of India Payment Data.

Figure 2: Share of amount transacted using mobile payments by households



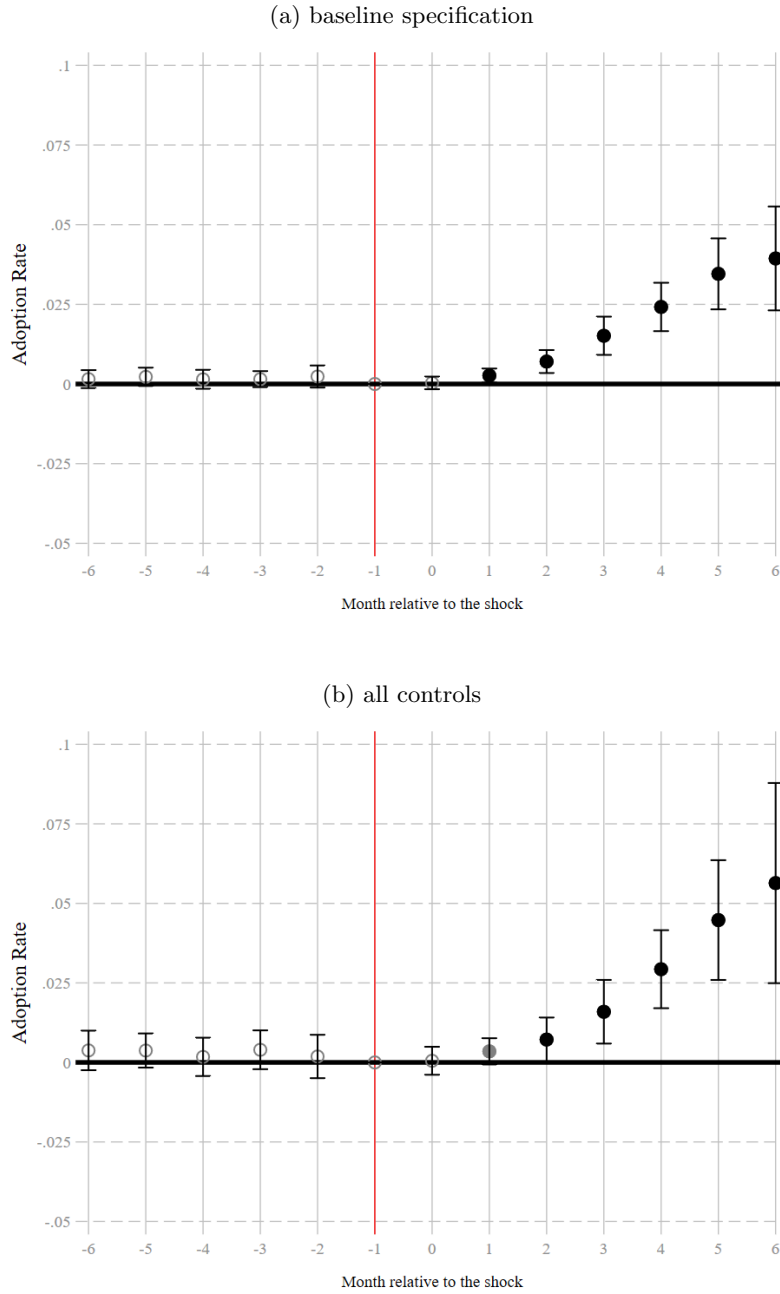
Notes: The figure plots the share of the amount transacted by households using mobile payments in the bank-level transaction data. All figures report the average of the share across 20 age bins and the slope of the line is reported in each panel. Panel (a) reports the mean with no controls. Panel (b) reports the means with baseline demographic controls for gender, marital status, and occupation. Panel (c) reports the means with baseline controls as well as pincode-month controls. Panel (d) reports the means with baseline controls as well as pincode-wealth bins-month controls. Panel (e) reports the means based on Panel (d) but is conditional on the sample of households that also hold a credit card. Each figure also reports the estimated coefficient β from the regression of the share of mobile payments on age with the controls based on the corresponding figure.

Figure 3: District Adoption Dynamics
 (New Store Adopting/Total firms per district ('100))




Notes: The figure plots the dynamic treatment effects of age structure on adoption. The dependent variable is the number of stores that adopted our fintech company in month t and district d , scaled by the total number of firms in the districts (in hundreds) measured by the Census. The period considered is the six months before and after May 2019 (i.e., zero in the graph), which is the month when the company introduced the mobile payment option. The graphs report the coefficients β_k from specification 5. Panel (a) reports the effects from baseline specification without any baseline district-level controls; panel (b) reports the effects from the specification that includes the district controls interacted with month fixed effects. Baseline district controls include the population (IHS), the share of agricultural workers, the number of firms (IHS), literacy rate, the share of working population, and the log of average night lights in 2018 in the district. Vertical lines indicate 95% confidence intervals. Black dots represent significance at 95% significance levels, gray dots represent significance at 90% significance levels, and hollow dots represent insignificant levels. Standard errors are clustered at the district level.

Figure 4: District Adoption Dynamics: IV results



Notes: The figure plots the dynamic treatment effects of age structure on adoption, where the share of young adults is instrumented by a function of the historical sex-ratio, as described in the paper. The dependent variable is the number of stores that adopted our fintech company in month t and district d , scaled by the total number of firms in the districts (in hundreds) measured by the Census. The period considered is the six months before and after May 2019 (i.e., zero in the graph), which is the month when the company introduced the mobile payment option. The graphs report the coefficients β_k from specification 5. Panel (a) reports the effects from baseline specification without any baseline district-level controls; panel (b) reports the effects from the specification that includes the district controls interacted with month fixed effects. Baseline district controls include the population (IHS), the share of agricultural workers, the number of firms (IHS), literacy rate, the share of working population, and the log of average night lights in 2018 in the district. Vertical lines indicate 95% confidence intervals. Black dots represent significance at 95% significance levels, gray dots represent significance at 90% significance levels, and hollow dots represent insignificant levels. Standard errors are clustered at the district level.

Figure 5: Adoption across pincodes: university areas



Notes: The figure plots the dynamic treatment effects of the presence of a university on the adoption of our fintech company. The dependent variable is the (inverse hyperbolic sine transformation of) the number of stores that adopted our fintech company at pincode-level in a month. The graphs report the coefficients γ_k from specification 6, and always include district-by-month fixed effects as well as pincode fixed-effects. The period considered is the six months before and after May 2019 (i.e., zero in the graph), which is the month when the company introduced the mobile payment option. Vertical lines indicate 95% confidence intervals. Black dots represent significance at 95% significance levels, gray dots represent significance at 90% significance levels, and hollow dots represent insignificance at 90% levels. Standard errors are reported in parentheses and are clustered at the pincode level.

Table 1: Variance composition

	Share of amount transacted with mobile money				
	(1)	(2)	(3)	(4)	(5)
Age	99%	81%	74%	65%	38%
Gender	1%	1%	1%	1%	0%
1(Married)		18%	16%	14%	8%
Industry			9%	8%	5%
Wealth				12%	7%
Pincode					42%

Notes: The table reports the variance decomposition generated using the Shapley Value approach as in (Huettnner and Sunder 2012). Each column reports the share of the outcome’s explained variance that is due to each of the characteristics reported across rows. Each characteristic is classified as a group rather than a continuous variable (for e.g., variable Age represents 48 bins, each corresponding to one age group between integer age 18 to 65), and the number reported represents the share explained by the whole group. Each column should sum to 100.

Table 2: Age Structure and Mobile Demand

	Adoption Rate		# Adoptions (IHS)	
	(1)	(2)	(3)	(4)
AgeStructure _d × Post _t	0.005*** [0.002]	0.008*** [0.002]	0.083*** [0.028]	0.162*** [0.036]
Observations	7,722	7,722	7,722	7,722
R-squared	0.559	0.602	0.839	0.849
District f.e.	✓	✓	✓	✓
Month f.e.	✓	✓	✓	✓
District Controls × Month f.e.	✗	✓	✗	✓

Notes: The table reports the difference-in-differences estimates of the effect of the age structure on adoption. The estimated specification is equation 4. Columns 1 - 2 report the estimate, where the outcome is expressed as the number of stores that adopted our fintech company in month t and district d , scaled by the total number of firms in the districts (in hundreds) measured by the Census. Columns 3 - 4 report the estimate on the IHS of the number of stores that adopted our fintech company district d during month t . Odd columns have no controls while even columns incorporate district controls interacted with month dummies. The district controls include the population (IHS), the share of agricultural workers, the number of firms (IHS), literacy rate, the share of working population, and the log of average night lights in 2018 in the district. Standard errors are reported in parentheses and are clustered at the district level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3: Age Structure: Balance Table

	Univariate OLS		Baseline Controls	
	Coefficient (1)	R^2 (2)	Coefficient (3)	R^2 (4)
Population (IHS)	-0.192*** (0.045)	0.042		
Share of agricultural workers	-0.023*** (0.005)	0.048		
Number of firms (IHS)	-0.360*** (0.047)	0.116		
Literacy Rate	-0.049*** (0.004)	0.222		
Share of working population	-0.015*** (0.003)	0.051		
Night lights (IHS)	-0.069 (0.063)	0.003		
Total stores (IHS)	-0.610*** (0.080)	0.077	-0.056 (0.082)	0.612
Total new stores (IHS)	-0.512*** (0.070)	0.065	-0.118 (0.072)	0.562
Total transaction volume (IHS)	-0.784*** (0.105)	0.081	-0.055 (0.121)	0.548
Total transaction amount (IHS)	-0.859*** (0.144)	0.055	-0.018 (0.189)	0.417
Total rural population (IHS)	0.003 (0.105)	0.000	0.046 (0.119)	0.270
Number of schools (IHS)	-0.120*** (0.044)	0.023	0.039 (0.035)	0.781
Population density	-0.127*** (0.048)	0.017	-0.075 (0.066)	0.488
Bank Branch Density	-0.113* (0.059)	0.005	-0.010 (0.064)	0.044
Share of manufacturing workers	-0.005 (0.005)	0.002	-0.003 (0.008)	0.221
Share of small firms	-0.000 (0.000)	0.000	-0.000 (0.001)	0.257
Share of primary education	-0.049*** (0.007)	0.112	-0.005 (0.004)	0.808
% of urban HH with mobile phones	-0.005 (0.006)	0.001	-0.013 (0.009)	0.053
% of urban HH with computers	0.002 (0.004)	0.001	0.006 (0.004)	0.356

Notes: The table tests for differences in observable district characteristics and age structure of the districts. Column 1 reports the mean of the district characteristics. The treatment variable is our measure of AgeStructure_d , as described Section 4. Columns 2 and 3 report the coefficient of the univariate OLS regression of each variable on the treatment variable. Columns 4 and 5 report the coefficients after controlling for the districts' population (IHS), the share of agricultural workers, the number of firms (IHS), literacy rate, the share of working population, and the log of average night lights in 2018 in the district. The district characteristics come from the 2011 Census, with the exception of the night light (which comes from the VIIRS Night light data) and information about the number of stores and transactions, which are measured using the data from our fintech company in the standard pre-period of the analysis. Robust standard errors are in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4: Age Structure and Adoption: IV analysis

	First Stage	2SLS	
	AgeStructure _d × Post _t (1)	Adoption rate (2)	# Adoptions (IHS) (3)
(Sex Ratio) _{d,1991} × Post _t	61.04*** (11.71)		
(Sex Ratio) ² _{d,1991} × Post _t	-25.70*** (5.332)		
AgeStructure _d × Post _t		0.020*** (0.0046)	0.256*** (0.085)
Observations	7,722	7,722	7,722
SW <i>F</i> -statistic	43.46		
District f.e.	✓	✓	✓
Month f.e.	✓	✓	✓
Controls × Month f.e.	✓	✓	✓

Notes: The table reports the instrumental variables (IV) estimates of the effect of the age structure on adoption. The estimated specification is equation 4, where we instrument the age structure of the district using a quadratic polynomial of sex ratio. Column 1 reports the first stage estimates. Column 2 reports the IV-2SLS estimate on our standard outcome (i.e., number of new stores adopting in month t and district d , divided by the number of firms in the district, in hundreds). Column 3 reports the IV-2SLS estimate on the IHS of the number of new firms that obtained a terminal from the firm. District controls include the population (IHS), the share of agricultural workers, the number of firms (IHS), literacy rate, the share of working population, and the log of average night lights in 2018 in the district. Standard errors are reported in parentheses and are clustered at the district level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5: Adoption and University

	# Adoptions (IHS)		# Adoptions	
	(1)	(2)	(3)	(4)
$1(\text{has university})_p \times \text{Post}_t$	0.232*** [0.051]	0.128*** [0.030]	3.987*** [1.112]	2.832*** [0.712]
Observations	109,058	108,402	109,058	108,402
R-squared	0.695	0.762	0.512	0.605
Pincode f.e.	✓	✓	✓	✓
Month f.e.	✓	✓	✓	✓
District f.e. × Month f.e.	✗	✓	✗	✓

Notes: The table reports the difference-in-differences estimates of the effect of the presence of a university on the demand for mobile payments by retailers. The estimated specification is equation 6. Columns 1 - 2 report the estimate on the IHS of the number of new stores that adopted a terminal from our fintech company in pincode p during month t . Columns 3 - 4 report the estimate of the (raw) number of new stores that adopted a terminal from our fintech company in pincode p during month t . All columns include pincode fixed effects, month fixed effects and district-month fixed effects. Standard errors are reported in parentheses and are clustered at the pincode level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Internet Appendix for “Demographics and
Technology Diffusion: Evidence from Mobile
Payments”

A.1 Proofs

Proof of Lemma 1. Let $b(j) = \ln(a(j))$. We have:

$$\begin{aligned} s_Y(p(j), a(j)) &= \mathbb{P}\left(j = \arg \max_{k \in [1, \dots, J]} -\frac{\ln(p(k))}{\nu - 1} + \varepsilon(i, k)\right) \\ &= \int_{\varepsilon \in \mathbb{R}} f(\varepsilon; b(j)) \prod_{k \neq j} \mathbb{P}\left(\varepsilon(i, k) \leq \varepsilon - \ln\left(\left[\frac{p(j)}{p(k)}\right]^{\frac{1}{\nu-1}}\right)\right) d\varepsilon \end{aligned} \quad (\text{A1})$$

Given the assumptions on the distribution of shocks:

$$\begin{aligned} f(\varepsilon; b(j)) &= \exp(-(\varepsilon - b(j)) - \exp(-(\varepsilon - b(j)))) \\ \mathbb{P}\left(\varepsilon(i, k) \leq \varepsilon - \ln\left(\left[\frac{p(j)}{p(k)}\right]^{\frac{1}{\nu-1}}\right)\right) &= \exp\left[-\left[\frac{p(j)}{p(k)}\right]^{\frac{1}{\nu-1}} \exp(-(\varepsilon - b(k)))\right] \end{aligned}$$

With these expressions, the integrand in Equation (A1) can be written as:

$$\exp\left[-(\varepsilon - b(j)) - (1 + \Sigma(j)) \exp(-(\varepsilon - b(j)))\right], \quad (\text{A2})$$

where:

$$\Sigma(j) \equiv \sum_{k \neq j} \frac{\exp(b(k))}{\exp(b(j))} \left[\frac{p(k)}{p(j)}\right]^{-\frac{1}{\nu-1}} \quad (\text{A3})$$

A primitive for the function in Equation (A2) is:

$$\frac{1}{1 + \Sigma(j)} \exp\left[-(1 + \Sigma(j)) \exp(-(\varepsilon - b(j)))\right].$$

Integration of Equation (A1) then gives the result. The derivations for the probability of an old consumer picking business j are similar, replacing $a(j)$ by 1. ■

Proof of Result 2. With this functional form for $c(a)$, the unique optimal technology adoption a is given by:

$$a = \frac{1}{2} + \sqrt{\frac{1}{4} + \frac{1}{\gamma J} \left(1 - \frac{1}{\nu}\right)}.$$

Therefore,

$$\frac{da}{d\gamma} = \frac{-\gamma^2 \theta}{2J} \left(1 - \frac{1}{\nu}\right) \left(\frac{1}{4} + \frac{\theta\gamma}{J} \left(1 - \frac{1}{\nu}\right)\right)^{-\frac{1}{2}} < 0,$$

$$\frac{d^2a}{d\gamma d\theta} = \frac{1}{2J} \left(1 - \frac{1}{\nu}\right) \left(\frac{1}{4} + \frac{\theta\gamma}{J} \left(1 - \frac{1}{\nu}\right)\right)^{-\frac{1}{2}} - \frac{1}{4J^2} \left(1 - \frac{1}{\nu}\right)^2 \left(\frac{1}{4} + \frac{\theta\gamma}{J} \left(1 - \frac{1}{\nu}\right)\right)^{-\frac{3}{2}}.$$

A necessary and sufficient condition for the latter expression to be positive is:

$$\frac{1}{4} + \frac{1}{2} \frac{\theta\gamma}{J} \left(1 - \frac{1}{\nu}\right) > 0,$$

which is always true. ■

A.2 Data Appendix

In this section, we discuss more in details some aspects of our data construction.

A.2.1 Bank Data

The first data used in the paper is a data set provided by one of the leading bank in India. As explained in Section 2, the data is a sample of about 200,000 customers from this bank, which is active across the whole country and several business areas. In this Appendix, we aim to expand a bit on some of the tests presented in the draft.

As part of the data validation conducted in the data section, we compare how our measure of age and total deposits in our bank data compares with representative data sets about Indian households. These analyses are conducted on the subset of individuals 18 to 65 years old, since this is the population later used in the analysis. In the bank data, the age of the account owner is provided as of January 2020. Furthermore, we estimate deposits from the bank data as the total deposit available in January 2020. The analyses use data from January and February 2020, the closest months to the fintech experiment that we were able to obtain from the bank.

To benchmark age, we use the NFHS survey conducted from 2019-2021, a nationally representative household level survey on household level demographics and health outcomes. The data set provides directly the age of household's head at the time of the survey, and this variable is used directly to construct our age distribution. To make the data representative, we employ the weights provided in the data set.⁶¹ We focus on household head because we want something that is comparable to the bank data. As illustrated in panel (a) of Figure A-1, the age profiles of bank account owners and household heads closely align, with a minor under-representation of individuals aged 60 to 65 offset by a higher presence of middle-aged individuals (30-50 years).

⁶¹https://dhsprogram.com/data/dataset/India_Standard-DHS_2020.cfm?flag=0

Despite a similar age profile, we expect our data to over-represent wealthier individuals. Mechanically, individuals in our data have a positive bank deposit balance. To benchmark the deposit distribution, we utilize the AIDIS survey, part of the 77th round of the NSS survey, using data conducted for the first visit in 2019.⁶² This is a nationally representative survey on households, reporting information about families' assets and liabilities. From this data, we obtain the value of deposits as of 06/30/2018 of households from the table called Visit1 Level - 12 (Block 11a) - Financial assets including receivables (other than shares and related instruments) owned by the household using assets with serial numbers 3-9 based on AIDIS survey 2019 documentation. In particular, these categories are: (a) 3 - deposit in savings bank account (excl. Post Office Savings Bank POSB); (b) 4 - fixed deposit/term deposit/ RD / flexi-RD in banks (excl. POSB); (c) 5 - savings and/or fixed deposits in post office savings bank; (d) 6 - other fixed income deposits (NSC, KVP, saving bonds, other small savings schemes, etc.); (e) 7 - deposits in cooperative banks; (f) 8 - deposits with non-banking finance companies; (g) 9 - deposits with Co-op credit society/micro-finance institutions/self-help groups. If households do not report assets with such serial numbers, we assume that they have 0 deposits.

Based on this, we obtain that 52% of the population does not have a bank deposit account. The final figure is then constructed conditional on the household having non-zero deposits, so that we can make this more directly comparable to the bank data. Also in this case, we use the internally provided survey weights to make the data representative. Appendix Figure A-1, panel (b), then confirms that our data over-sample households with higher deposit account balances.

While we recognize the difference in wealth between our sample and the population at large, we do not think that this difference hinders us to conduct a insightful study of payment behavior. First, in general, wealthier individuals are more inclined towards electronic payments. Therefore, although the data may not perfectly represent the entire Indian population, it offers a useful snapshot of the subset more engaged with electronic payments. Second, even if the data set is skewed towards wealthier individuals, our dataset has a broad coverage across all wealth levels. This feature allows us to directly control for differences in wealth, and therefore isolate the effect of age from wealth differences.

A.2.2 Fintech Payment Data

We now describe in detail the way we construct the district panel measuring the growth of our fintech payment provider in India. The raw data is provided to us in the form of a transaction panel identified by a terminal and firm ID as well as a master file for the terminals that provides us with the pincode where the terminal operates, and some firm demographic information for each terminal. To have a sense of the data set, we have about 440M transactions, covering the payment activity of more than 900K firms. We want to note here that the data from the fintech company is completely orthogonal to the data from the bank, discussed in Section 2.

The data cleaning for the store panel involves a few key steps. First, we consolidate the terminal

⁶²See <https://microdata.gov.in/nada43/index.php/catalog/156/overview>

transaction panel and identify payment processor and transaction types. This data at the timestamp level is aggregated to a month-transaction category-payment processor panel. The second step is to essentially match the transaction file to the master file. This step allows us to attach information like the pincode. We want to note here that our data does not allow us to identify the firm and therefore we cannot match this data to any external data set at the firm level. There are two small issues with the master file that need to be solved. First, the original master file has issues due to changing terminal IDs over time by our fintech company. We solved this issue with the help of additional datasets provided by the fintech company that provides information on how the IDs change. Second, the pincode information is available for only around 90% of terminals. For the rest, we have a location ID variable, which is an internal ID variable constructed by the company. The good news is that this location ID uniquely identifies pincodes except for a handful of terminals (i.e., seven terminals across the whole data set). We then use this variable to fill in remaining pincodes. When there are more than one location IDs across the files provided by the company, we use the one in the most recent file.

In the last step, we match the two data sets. In a few cases, we do not find all terminals in the constructed master file (about 10% of the sample). However, since both datasets have a firm identifier, we can infer the location for some of these unmatched terminals. In particular, a sizable subset of these firms only operates terminals in one pincode, and therefore we infer that the unmatched terminal is likely also in the same pincode. We discard terminals with more irregular patterns. However, as it will be clear below, this choice is not going to affect our analysis unless the specific discard terminal is the first terminal adopted by the firm.

Once we have connected to each terminal the firm ID and the location, we construct our store ID: as we mentioned in the paper, we define a store as a combination of terminals belonging to the same firm within the same (6-digit) pincode. In general, the largest majority of firms only have one terminal. The construction of the store ID allows us to normalize for the fact that certain stores may have more than one terminal, and in some cases firms have more than one location. However, it is important to point out that this adjustment is likely second order here, since the majority of firms have only one terminal, and only a few thousand firms have more than one location. This is a relatively small number given that the total size of the data.

We then use the information about transactions to understand the adoption time, which we infer by looking at the time of the first transactions done in a store. Notice that for a subset of the sample, we also have information about the time of installation of the terminal, which can be used as an alternative way to measure the exact time of adoption. Comparing adoption month using the first transaction and installation for the sample of terminals adopted during the sample period, we find that the two measures coincide exactly almost 86% of the sample, and the gap is limited to a few days for most others (for instance, if we allow the first transaction to be a month delay, the measures are the same for over 94%). We therefore use the first day of adoption as our main measure because this version is available for our full sample, rather than a portion of it.

The last step is straightforward: once the data is organized, we aggregate the data at district

by month level, measuring the number of stores active in a district as well as the number of new stores adopting our company’s payment that specific district. We particularly focus our analysis on the period around the May 2019 policy shift, which is what we use in our model.

A.2.3 Data on University Location

This section outlines details regarding the presence of universities at the pincode level, as discussed in the main text. The goal of utilizing this data is to pinpoint locations with an unusually high concentration of young adults. Specifically, the aim is to compile an exhaustive list of higher education institutions operational in 2019, the year under scrutiny in our analysis. Collecting this data presented three main challenges, which are outlined below. Firstly, we needed to secure a reliable list of higher education institutions. Secondly, it was important to verify ”to the best of our ability” that these institutions were indeed active in 2019. Lastly, identifying the pincode for each institute was necessary.

To determine the list of universities in India, we utilize the classification of universities provided by the University Grants Commission of India, an organization that provides recognition to universities in India. We utilize four groups of universities provided by the UGC:

1. Central Universities: established by an act of parliament and are under the purview of the Department of Higher Education in the Ministry of Education;
2. State Universities: established by an act of parliament and are under the purview of the Department of Higher Education in the Ministry of Education;
3. Deemed Universities: status of autonomy granted by the Department of Higher Education on the advice of the UGC, under Section 3 of the UGC Act;
4. Private Universities: approved by the UGC. They can grant degrees, but they are not allowed to have off-campus affiliated colleges.

In addition to this list, we also use the list of Institutions of National Importance, which are not universities but considered important by the Indian Ministry of Education.

We now provide a bit more detail about each of these lists, in particular describing how we make sure that these centers were active in 2019. Regarding Central Universities, the list was provided by UGC, with a document that appeared to be released in April 2023.⁶³ One concern is that some Central Universities may have been added after 2020: we then manually check if there was any law passed in 2020 about this and we could not find any. For State Universities, the list was also provided by UGC.⁶⁴ In this case, the last was provided as of March 31st 2019, therefore matching perfectly our time of interest. For Deemed Universities, the list was provided on the UGC

⁶³https://web.archive.org/web/20230404082827/https://www.ugc.gov.in/oldpdf/Consolidated_CENTRAL_UNIVERSITIES_List.pdf

⁶⁴https://web.archive.org/web/20190805020657if_/https://www.ugc.ac.in/oldpdf/State%20University/Consolidated%20State%20%20University%20List.pdf

website,⁶⁵ updated at the day close to the download (i.e., January 2024). In this case, it could therefore be possible that some universities in the list were opened after 2020: we manually checked a sub-sample of the data and could not find any cases. Therefore, even if we cannot exclude this issue entirely, we do not expect this problem to be significant. The list of private Universities was also found on an external website (i.e., Boston University) but appeared to come from UGC, and the list was updated November 12th 2018, therefore fitting well our needs.⁶⁶ Lastly, the list of Institutions of National Importance was found on the Government website updated at least until 2022.⁶⁷ A manual check of the list seems to exclude any recent additions.

We then combine the list of universities, cleaned the data, also removing a few duplicates that are found in the process. We then connect each university with a (six-digit) pincode: for entries coming from UGC files, the pincode can be generally extracted from the address that is provided. For Institutions of National Importance, the address is not provided and we had to manually add the pincode. In general, we add pincode manually using Google, specifically searching "university name" " "pincode". Two points related to this analysis are worth highlighting. First, it is important to point out that the pincode identified through this data collection is likely to capture the location of the headquarters or main building of the University. For Universities that are very large, it is possible that some buildings are located outside the original pincode. Given the impossibility to find a complete list of all Universities' buildings in India, we thought that this issue was generally acceptable. Furthermore, we expect that, if anything, this issue would bias us towards finding no difference in the data. This is particularly the case given that this analysis will only exploit within-district variation.

Second, the higher education sector in India is characterized by the presence of both universities and colleges. Universities, which we consider above, are educational institutions that are authorized to award degree under a Central or a state Legislature. In other words, these intuitions are close to what is traditionally referred to as a university or 4-year college in the U.S.. However, India is also characterized by another type of higher education institution, generally referred to as Colleges. These institutions are not authorized to award an educational degree in their own name and may be affiliated with some university. However, colleges are much smaller than universities in their enrollment size: two-third of colleges have less than 500 students and only 8% colleges have greater than 2000 students.⁶⁸ Moreover, 60% of the colleges are located in rural areas. Thus, unlike universities, we do not expect the presence of colleges to be a large enough force to have significant impact on local demand coming from younger customers (i.e. college graduates). Furthermore, their presence is likely to also bias our findings towards zero.

⁶⁵<https://deemed.ugc.ac.in/Home/ListOfDeemedToBeUniversity>

⁶⁶<https://www.bu.edu/globalprograms/files/2019/02/Private-University-Consolidated-List-Private-Universities.pdf>

⁶⁷<https://www.education.gov.in/institutions-national-importance>

⁶⁸See the All India Survey of Higher Education (2018) from the Ministry of Education India: <https://cdnbbsr.s3waas.gov.in/s392049debbe566ca5782a3045cf300a3c/uploads/2024/02/2024021480881112.pdf>


Appendix Figures and Tables

Figure A-1: Comparison of distributions across datasets



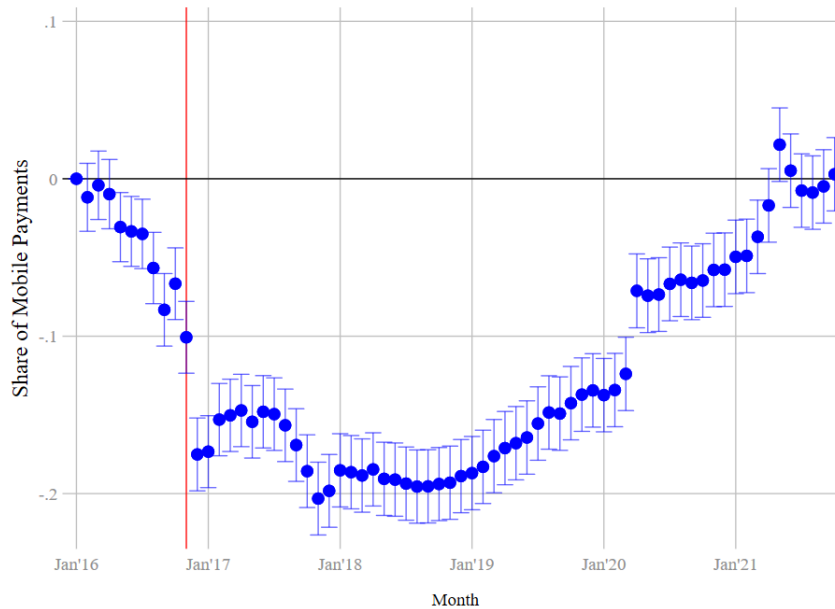
Notes: This figure compares the age and income distribution in our bank level data to the same information provided in nationally representative surveys. Panel (a) reports the age distribution, dividing the sample in 5-year intervals (with the exception of the youngest group that goes from 18 to 25). For each group, the first bar reports the share of head of the household in that age group from the NFHS 2019-2021 survey, as described in the paper; the second bar reports the same statistic for our bank data. Panel (b) reports the wealth distribution, across three broad category (i.e., less than 5,000 Rp., between 5,000 and 100,000, and above 100,000). For each group, the first bar reports the share of individuals that have that level of deposit in the AIDIS 2019, as described in the paper; the second bar reports the same statistic for our bank data. Notice that, as explained in the paper, the share from the AIDIS is conditional on having any deposit.

Figure A-2: Share of amount transacted using mobile payments by households



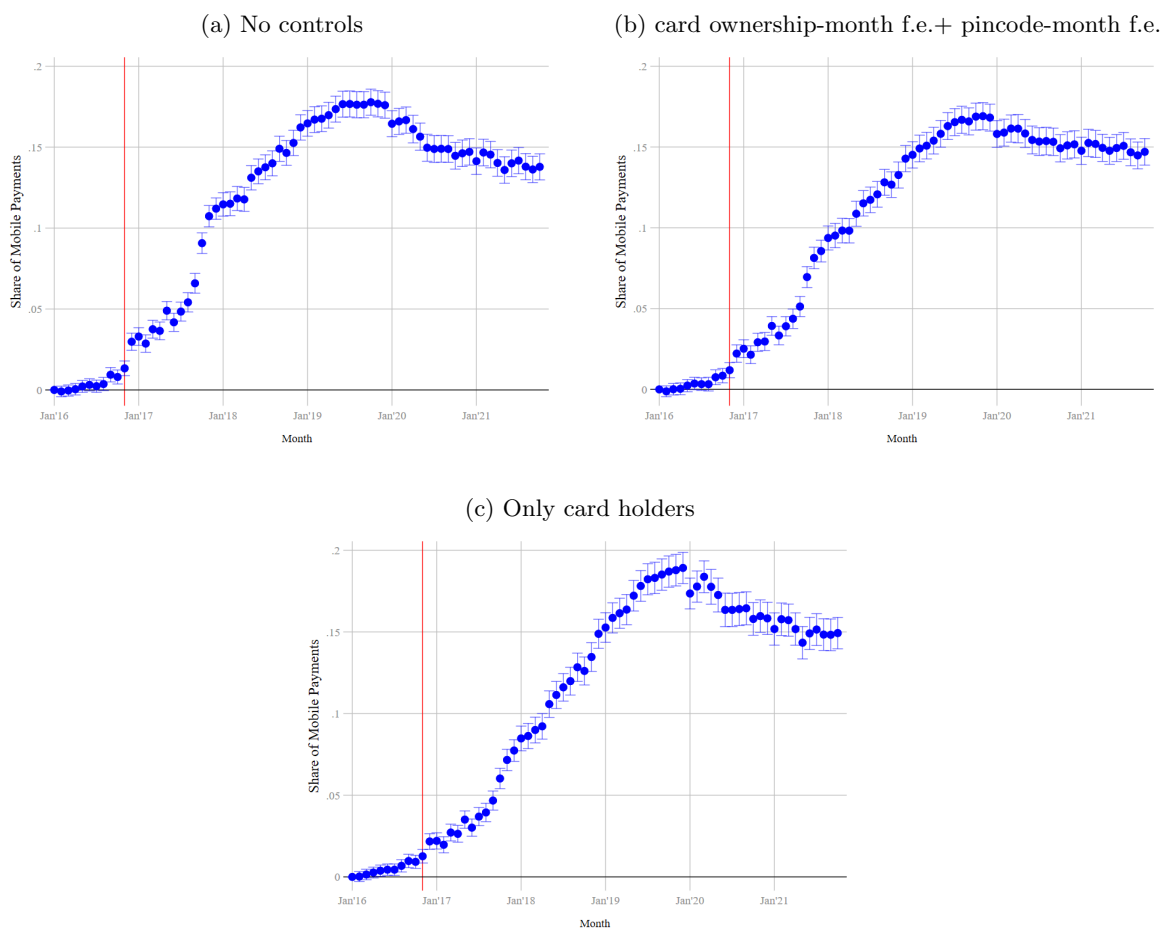
Notes: The figure plots the estimates of the share of the amount transacted using mobile payments by households across different age-groups. We normalize the age groups of 18 to 24 to be zero. 95% confidence intervals are denoted using the blue shaded region.

Figure A-3: Difference in Mobile Penetration and Card Early-Users



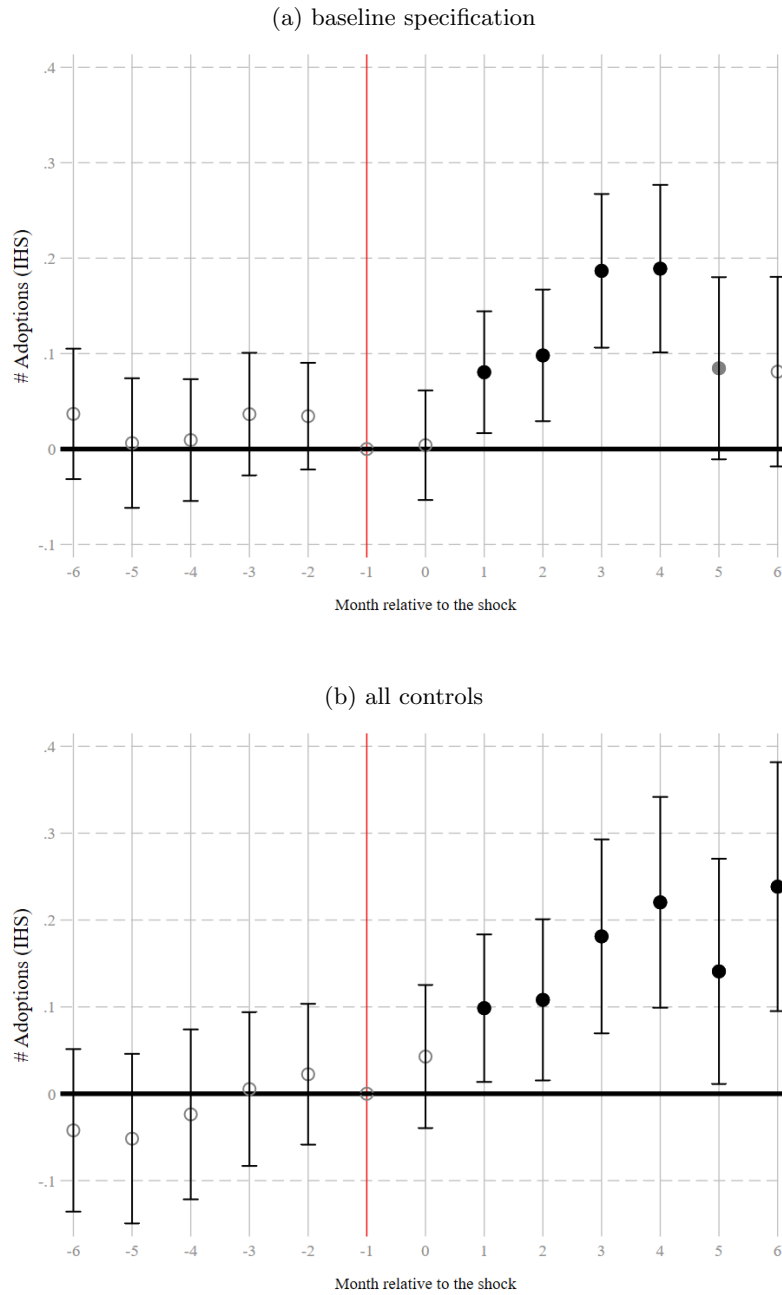
Notes: The figure plots month-by-month the difference in the share of mobile transactions between consumers that were early users of credit cards (i.e., had credit card in 2015) and not early users. The share of mobile is defined as usual. The vertical red-line identifies November 2016 — the month of Indian Demonetization. The sample uses the set of customers active between (at least) 2015 and 2021. The point estimate is also reported together with the 95% confidence interval.

Figure A-4: Difference in Mobile Penetration across Young and Old



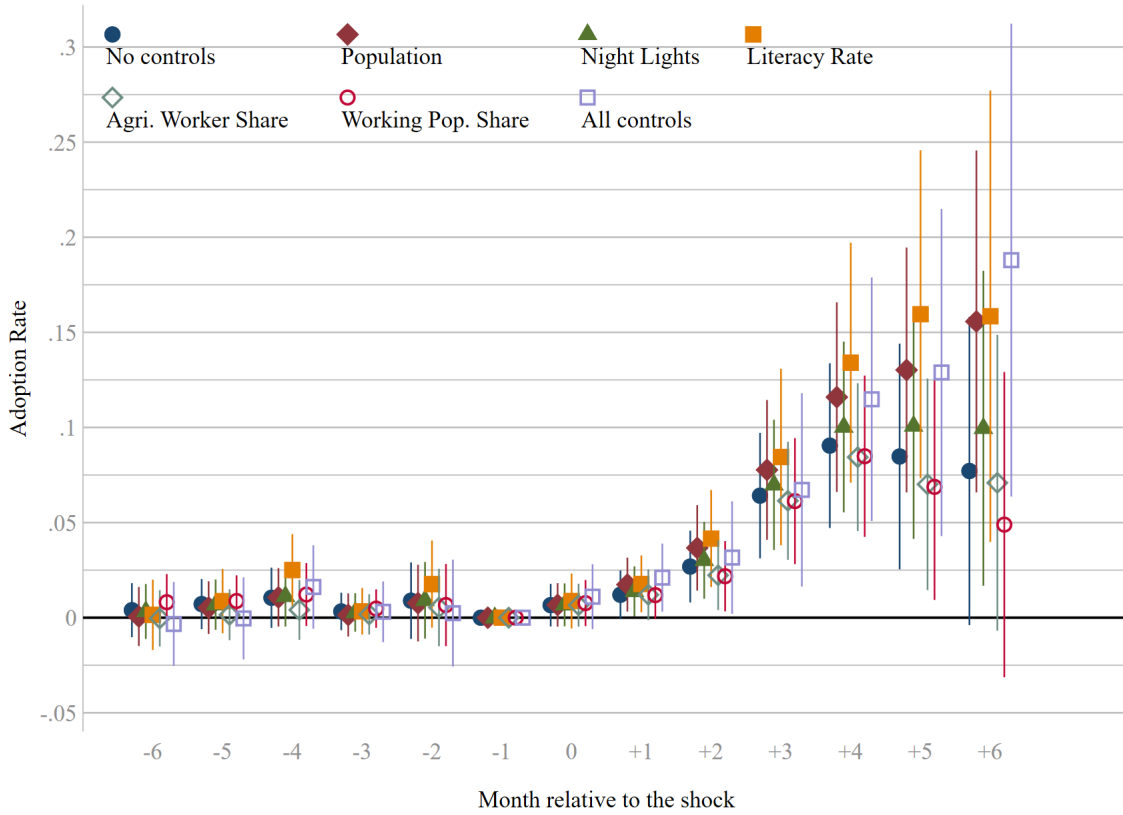
Notes: The figure plots month-by-month the difference in the share of mobile transactions between young and old consumers using the panel from the the bank-level transaction data. We divide the sample across young and old splitting around age 40, defined in 2015. The share of mobile is defined as usual. Panel (a) reports the difference with no controls. Panel (b) reports the difference while also controlling for card ownership-month fixed-effects and pincodes-month fixed-effects. Panel (c) reports the difference using only the sample of those borrowers that have used any card in the sample period. The vertical red-line identifies November 2016 — the month of Indian Demonetization. The sample uses the set of customers active between (at least) 2015 and 2021. The point estimate is also reported together with the 95% confidence interval.

**Figure A-5: District Adoption Dynamics:
IHS of New Stores Adoption per district**




Notes: The figure plots the dynamic treatment effects of age structure on adoption. The dependent variable is the inverse hyperbolic sine (IHS) transformation of number of stores that adopted our fintech company in month t and district d . The period considered is the six months before and after May 2019 (i.e., zero in the graph), which is the month when the company introduced the mobile payment option. The graphs report the coefficients β_k from specification 5. Panel (a) reports the effects from baseline specification without any baseline district-level controls; panel (b) reports the effects from the specification that includes the district controls interacted with month fixed effects. Baseline district controls include the population (IHS), the share of agricultural workers, the number of firms (IHS), literacy rate, the share of working population, and the log of average night lights in 2018 in the district. Vertical lines indicate 95% confidence intervals. Black dots represent significance at 95% significance levels, gray dots represent significance at 90% significance levels, and hollow dots represent insignificant levels. Standard errors are clustered at the district level.

**Figure A-6: District Adoption Dynamics
Sensitivity to Controls**




Notes: This figure reports a robustness of our main specification to the inclusion of controls. In particular, we reproduce the same Figure 3 with different level of controls. As in the main figure, the dependent variable is the number of stores that adopted our fintech company in month t and district d , scaled by the total number of firms in the districts (in hundreds) measured by the Census. The period considered is the six months before and after May 2019 (i.e., zero in the graph), which is the month when the company introduced the mobile payment option. However, each set of coefficient differs in the controls used: in particular, we consider the specification without any control (as in panel a of Figure 3) as well as with each control included alone. As a benchmark, we also report the specification with all the controls (as in panel b of Figure 3) Vertical lines indicate 95% confidence intervals. Standard errors are clustered at the district level.

**Figure A-7: District Adoption Dynamics:
no adult adjustment**




Notes: This figure provides a robustness test for the main dynamic specification. Everything is identical to the main figure (i.e., Figure 3), but for the treatment variable. In the main analysis, the treatment variable is the number of individuals between 15 and 29, scaled by the number of adults, defined as individuals between 15 and 74. This robustness figure instead uses as treatment a measure that scales number of individuals between 15 and 29 by total population, without any adjustment for children or elderly. Apart from this change, everything is equivalent to the specification with controls. In particular, the dependent variable is the number of stores that adopted our fintech company in month t and district d , scaled by the total number of firms in the districts (in hundreds) measured by the Census. The period considered is the six months before and after May 2019 (i.e., zero in the graph), which is the month when the company introduced the mobile payment option. The graphs report the coefficients β_k from specification 5. Baseline district controls include the population (IHS), the share of agricultural workers, the number of firms (IHS), literacy rate, the share of working population, and the log of average night lights in 2018 in the district. Vertical lines indicate 95% confidence intervals. Black dots represent significance at 95% significance levels, gray dots represent significance at 90% significance levels, and hollow dots represent insignificant levels. Standard errors are clustered at the district level.

**Figure A-8: District Adoption Dynamics:
alternative treatment**




Notes: This figure provides a robustness test for the main dynamic specification. Everything is identical to the main figure (i.e., Figure 3), but for the treatment variable. In the main analysis, the treatment variable is the number of individuals between 15 and 29, scaled by the number of adults. This robustness figure instead uses as a treatment the share of individual that are less than 40. Apart from this change, everything is equivalent to the specification with controls. In particular, the dependent variable is the number of stores that adopted our fintech company in month t and district d , scaled by the total number of firms in the districts (in hundreds) measured by the Census. The period considered is the six months before and after May 2019 (i.e., zero in the graph), which is the month when the company introduced the mobile payment option. The graphs report the coefficients β_k from specification 5. Baseline district controls include the population (IHS), the share of agricultural workers, the number of firms (IHS), literacy rate, the share of working population, and the log of average night lights in 2018 in the district. Vertical lines indicate 95% confidence intervals. Black dots represent significance at 95% significance levels, gray dots represent significance at 90% significance levels, and hollow dots represent insignificant levels. Standard errors are clustered at the district level.

Figure A-9: District Adoption Dynamics
 (New Adopters/Population per district ('100,000))




Notes: This figure provides a robustness test for the main dynamic specification. Everything is identical to the main figure (i.e., Figure 3), but for the way the outcome is constructed. In particular, the dependent variable is the number of stores that adopted our fintech company in month t and district d , scaled by the total population (in hundred of thousands) measured by the Census. The period considered is the six months before and after May 2019 (i.e., zero in the graph), which is the month when the company introduced the mobile payment option. The graphs report the coefficients β_k from specification 5. Baseline district controls include the population (IHS), the share of agricultural workers, the number of firms (IHS), literacy rate, the share of working population, and the log of average night lights in 2018 in the district. Vertical lines indicate 95% confidence intervals. Black dots represent significance at 95% significance levels, gray dots represent significance at 90% significance levels, and hollow dots represent insignificant levels. Standard errors are clustered at the district level.

Figure A-10: Effect on Platform Size
 (# firms on platform / # firms per district ('100))




Notes: The figure plots the dynamic treatment effects of age structure on the total number of firms that are in our fintech platform. Apart from the outcome, everything is identical to the main figure (i.e., Figure 3). The dependent variable is the number of stores are in the platform in month t and district d , scaled by the total number of firms in the districts (in hundreds) measured by the Census. The period considered is the six months before and after May 2019 (i.e., zero in the graph), which is the month when the company introduced the mobile payment option. The graphs report the coefficients β_k from specification 5. Panel (a) reports the effects from baseline specification without any baseline district-level controls; panel (b) reports the effects from the specification that includes the district controls interacted with month fixed effects. Baseline district controls include the population (IHS), the share of agricultural workers, the number of firms (IHS), literacy rate, the share of working population, and the log of average night lights in 2018 in the district. Vertical lines indicate 95% confidence intervals. Black dots represent significance at 95% significance levels, gray dots represent significance at 90% significance levels, and hollow dots represent insignificant levels. Standard errors are clustered at the district level.

Figure A-11: Correlation: Sex Ratio (1991) and Age Structure (2011)




Notes: The figure documents the relationship between district sex ratios, defined as the number of males per female, in 1991 and the share of 2011 population in districts that is below 30 years. Each dot represents a district in the 2011 Census. The black line represents a quadratic polynomial fit.

**Figure A-12: Adoption across pincodes:
University areas, in levels**



Notes: The figure provides a robustness to the main university result (Figure 5). In particular, we change the way the outcome is measure. The dependent variable is the (raw, without any transformation) the number of stores that adopted our fintech company at pincode-level in a month. The graphs report the coefficients γ_k from specification 6, and always include district-by-month fixed effects as well as pincode fixed-effects. The period considered is the six months before and after May 2019 (i.e., zero in the graph), which is the month when the company introduced the mobile payment option. Vertical lines indicate 95% confidence intervals. Black dots represent significance at 95% significance levels, gray dots represent significance at 90% significance levels, and hollow dots represent insignificance at 90% levels. Standard errors are reported in parentheses and are clustered at the pincode level.


**Figure A-13: Adoption across pincodes:
University areas, only districts w. university**




Notes: The figure provides a robustness to the main university result (Figure 5). In particular, everything is identical to the main analysis, with the exception that here we only use the sample of districts that have at least one university in their territory. The dependent variable is the (inverse hyperbolic sine transformation of) the number of stores that adopted our fintech company at pincode-level in a month. The graphs report the coefficients γ_k from specification 6, and always include district-by-month fixed effects as well as pincode fixed-effects. The period considered is the six months before and after May 2019 (i.e., zero in the graph), which is the month when the company introduced the mobile payment option. Vertical lines indicate 95% confidence intervals. Black dots represent significance at 95% significance levels, gray dots represent significance at 90% significance levels, and hollow dots represent insignificance at 90% levels. Standard errors are reported in parentheses and are clustered at the pincode level.

**Figure A-14: Adoption across pincodes:
University areas, merchant-level variation**


Panel (a): Student-consumer businesses



Panel (b): Student-consumer businesses, broader




Panel (c): Placebo Merchants



Notes: The figure provides a robustness to the main university result (Figure 5). Across the three panels, we replicate our main analyses using only a sub-sample of merchants. In panel (a), we identify businesses that are mostly depending on local demand by students (i.e., retailers, gas stations, restaurants, leisure facilities, personal services, and transportation). Panel (b) modifies this definition to include also financial services, healthcare (e.g., pharmacies), and educational services. Panel (c) instead runs the analyses on businesses that should not be affected by the local demand by students, such as government and regulated sectors, manufacturing, wholesale, warehouse operations, and professional services. Across the panels, the dependent variable is the (inverse hyperbolic sine transformation of) the number of stores that adopted our fintech company at pincode-level in a month for the subset of merchants considered. The graphs report the coefficients γ_k from specification 6, and always include district-by-month fixed effects as well as pincode fixed-effects. The period considered is the six months before and after May 2019 (i.e., zero in the graph), which is the month when the company introduced the mobile payment option. Vertical lines indicate 95% confidence intervals. Black dots represent significance at 95% significance levels, gray dots represent significance at 90% significance levels, and hollow dots represent insignificance at 90% levels. Standard errors are reported in parentheses and are clustered at the pincode level.

Figure A-15: Pix Adoption in Brazil and Share of Young Individuals



Notes: The figure examines the relationship between the use of Pix in Brazilian municipalities and the age structure of the municipalities, following a similar graphical representation presented before. The use of Pix is measured as the average Pix used (in value) per unit of population in Brazil. Pix is measured using monthly data between November 2020 and July 2024. The age structure is measured using the share of population that is below forty years old, from the 2010 Census. The figure reports the scatterplot of the two quantities after controls and the linear fit. The slope coefficient of the linear fit is also reported. We control for (log) population, average income, literacy rate, and share of rural population.

Table A-1: Age Structure and share of new stores that adopted QR code

	% new stores that adopted QR code	
	(1)	(2)
AgeStructure _d	0.033*** (0.012)	0.048*** (0.015)
Observations	580	580
R-squared	0.012	0.139
District Controls	✗	✓

Notes: The table reports the estimates of the effect of the age structure in the district on the share of new adopters that adopted QR code with the company. The outcome is the share between the total number of stores that have adopted the product of our fintech company with at least one terminal enabled to use QR code, and the total number of adopters (without any requirement to have a QR code enabled terminal). The share is constructed over the full period May 2019 and November 2019. Column 1 reports the estimate without any controls and Column 2 reports the estimate after controlling for baseline district controls of the population (IHS), the share of agricultural workers, the number of firms (IHS), literacy rate, the share of working population, and the log of average night lights in 2018 in the district. Standard errors are reported in parentheses and are clustered at the district level. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table A-2: University Analyses: merchant-level

Panel A: # Adoption (IHS)				
	(1)	(2)	(3)	(4)
$1(\text{has university})_p \times \text{Post}_t$	0.065*** (3.30)	0.076*** (3.75)	0.000 (0.00)	0.171*** (6.93)
Pincode FE	Y	Y	Y	Y
District \times Month FE	Y	Y	Y	Y
Outcome	Student businesses	Student businesses (expanded)	Placebo	Others
Adj R-Sq	0.674	0.693	0.310	0.628
Obs	109,626	109,626	109,626	109,626
Panel B: # Adoption				
	(1)	(2)	(3)	(4)
$1(\text{has university})_p \times \text{Post}_t$	0.291*** (2.64)	0.332*** (2.65)	-0.004 (-0.24)	2.494*** (4.42)
Pincode FE	Y	Y	Y	Y
District \times Month FE	Y	Y	Y	Y
Outcome	Student businesses	Student businesses (expanded)	Placebo	Others
Adj R-Sq	0.674	0.693	0.310	0.628
Obs	109,626	109,626	109,626	109,626

Notes: The table reports the results of the university analysis, where we focus only on a sub-set of businesses across the various columns. In column (1), we focus on businesses that are mostly depending on local demand by students (i.e., retailers, gas stations, restaurants, leisure facilities, personal services, and transportation). In column (2) we expand this definition to include also financial services, healthcare (e.g., pharmacies), and educational services. In column (3) we run the analyses on businesses that should not be affected by the local demand by students, such as government and regulated sectors, manufacturing, wholesale, warehouse operations, and professional services. Column (4) focuses instead on the residual merchants, that do not belong to either the group defined in column (2) or (3) (this is mostly businesses categorized as "miscellaneous".) The dependent variable is the number of adoptions of our fintech solution: panel (a) uses this outcome after applying the inverse-hyperbolic transformation; panel (b) instead looks at the value in level. The specifications include fixed effects for pincode and district-by-month. The coefficients for the interaction between the presence of a university and the post period are reported. Standard errors are clustered at the pincode level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.