Demographics and Technology Diffusion: Evidence from Mobile Payments*

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

Abstract

Using evidence from the adoption of mobile payments in India, we show that demographic structure shapes the diffusion of new technologies. Younger adults are far more likely to use mobile payments than older consumers. These consumer patterns create stronger incentives for businesses facing younger customers to adopt new payment technologies. Using store-level data on merchant adoption, we exploit cross-district variation in age structure—both directly and through instruments based on historical fertility—to estimate the causal effect of demographics on business adoption decisions. The results show that regions with younger populations experience a larger merchant adoption of mobile payments. A model in which consumer attitudes toward technology differ across age groups implies that diffusion is inefficiently slow and that a positive short-run subsidy can restore efficiency. Aging thus slows the spread of financial innovation.

Keywords: Payments, Technology Diffusion, Fintech, Demographic Structure, UPI.
JEL Classification: O33, G23, L86, J11.

* *This version: October, 2025.* 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 Prettner (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

[1]Other 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.

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

relevance to other environments beyond India.

Our analysis proceeds in four 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, consistent with our simple model. Fourth, we use the model to study whether heterogeneity in consumer attitudes toward technology leads to inefficient adoption decisions by businesses in competitive equilibrium, relative to what a planner would choose. We show that, even without externalities, technology diffusion is inefficiently low in the model, and that a positive subsidy to adoption can restore efficiency in the short run.

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] It allows us to measure the proportion of electronic payments made using mobile technologies. We establish two main stylized facts. First, we show that age is a primary factor explaining the variation in the mobile payment use among consumers: 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 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

[3]The 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.

[4]The 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.

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

we interpret as mobile payments. They face customers of two potential types, young 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.

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.

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

[6] We 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 further support our interpretation, we introduce a novel approach that leverages the distribution of universities within districts. The idea is that businesses near universities naturally cater to a younger clientele (i.e., students), yet should be otherwise similar to businesses in other neighborhoods within the same district. Using university presence at the pincode level—the most granular location identifier in our data—and controlling for district-by-month fixed effects, we show that post-May 2019 adoption rates rose more sharply in university pincodes than in other pincodes within the same district. Apart from allowing us to control for time-varying district-level confounders, this approach has the added benefit of clearly identifying the consumer group driving demand for the technology: university students. This specificity enables us to predict which types of businesses are most affected. We find that the results are driven mainly by businesses serving university students. In contrast, merchants less likely to cater to students show no significant effects, effectively functioning as a placebo group.[8]

Having validated the basic predictions of our model, we use it to asses whether technology adoption decisions of businesses are efficient when different consumer groups (young and old, in the context we study) value the technology differently. A simple prior here might be that, absent externalities specific to the technology (say, network externalities, in the case of payment systems), the adoption level is efficient. We show that this is not the necessarily the case. In our model, businesses restrict quantities so as to be able to charge higher prices. Because technology adoption and output are complements, the quantity restriction also leads to an inefficiently low level of

[7]The 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.

[8]For 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.

adoption, even when the technology features no other externalities. Thus monopoly distortions cause under-investment in the technology. However, surprisingly, the optimal adoption subsidy is independent of both the shape of adoption costs and the fraction of consumers who value the technology (i.e, the young), so long as at least some of them do. Absent network externalities, the subsidy only depends elasticity of demand with respect to price (which is the same across young and old consumers) and the share of household expenditure on local businesses, and has an interpretation as a Pigouvian wedge that neutralizes the effect of the monopoly distortion on the adoption margin. When allowing for network externalities, the wedge increases, though it remains independent of demographics and adoption costs. Thus in general, except when no consumer values the technology, adoption should be subsidized to offset the under-investment problem, yet not necessarily in a way that depends on the specifics of how consumers value the technology, or how costly it is for businesses to adopt.[9]

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. Moreover, business adoption decisions may not be efficient when different customer groups value the technology differently, potentially calling for adoption subsidies. 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 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, allowing us to isolate the effects of demographics from 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 6 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

[9]These results are obtained in a version of the model with a fixed number of firms, consistent with the fact that our empirical analysis focuses on relatively on a relatively short window of time after the technology is made available to retail businesses. However, we also study the model with free-entry. We show that in this case, adoption is inefficiently low so long as the technology features network externalities, so that a subsidy is always optimal. However, different from the model with a fixed number of firms, the subsidy may increase with the share of consumers who value the technology. So, under free-entry, while it is always optimal to subsidize adoption, there are grounds to do so more aggressively in places where the demographic base is younger. Section 5 discusses the intuition behind this result, and contrasts it with the case of a fixed number of firms in more detail, highlighting in particular an additional source of inefficiency, the love-for-variety effect, which also interacts with adoption decisions.

influence firms' adoption decisions. We also examine the welfare implications of this distortion.

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.[10] 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).[11] 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.[12]

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

[10]For 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.

[11]This 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; Agarwal et al. 2024; Bian et al. 2023) and businesses (Agarwal et al. 2019).

[12]Our 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.

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 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.[13] While the UPI was formally introduced in 2016, UPI transactions remained small, compared to mobile wallet transactions, until the end of 2017.[14] However, the

[13] In 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.

[14] https://www.npci.org.in/what-we-do/upi/product-statistics

UPI's growth trajectory surpassed that of mobile wallets post-2017, reaching approximately 80% of mobile transactions by the end of 2021.[15]

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.[16] 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 process itself. As the term suggests, mobile transactions are executed using an app on a phone or similar digital device.[17] 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.[18] 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

[15] The 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.

[16] For 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.

[17] In 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).

[18] While 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.

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.[19] 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.[20]

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 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.[21] 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

[19]See 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/.

[20]In 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.

[21]To 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).

transactions.[22] 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.[23] 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.[24]

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).[25]

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.[26] 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%).[27] 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.

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

[23] We choose to compare our data's age distribution with that of household heads as this characteristic is more likely to 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.

[24] Appendix A.1.1 contains a more detailed discussion of the comparison of the age and wealth distribution of our data with nationally representative samples of households.

[25] Several 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 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.

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

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

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.[28]

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.[29]

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.[30] Although the magnitude of the relationship between age and mobile payment usage is somewhat diminished, consistent evidence of a significant of relationship between age and mobile payment usage remains compelling.[31]

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

[28] In 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.

[29] We 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 for each of the 20 bins.

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

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

cards, potentially underpinning the observed relationship.[32] 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).[33]

Discussion After highlighting the surge of mobile payments in India in recent years, this section showed 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 (Dynan 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

[32]It 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.

[33]Due 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.

bank.[34] 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.[35]

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

[34] While 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.

[35] Importantly, 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 by a direct early use of cards.

toward technology vary with age, as documented in Section 2. In Section 5, we will also use the model to evaluate the effects of adoption subsidies.

3.1 Description

Consumers There is a unit mass of consumers, indexed by $i \in [0,1]$. Each consumer has preferences over an outside good, $O(i)$, and an aggregate of varieties produced by businesses in the economy, $C(i)$, described by:

$$W(i) = \log\left(O(i)^{1-\alpha} C(i)^{\alpha}\right), \tag{1}$$

where $\alpha \in [0,1]$ governs the elasticity of substitution between $O(i)$ and $C(i)$. The outside good serves as the numéraire, and each household is endowed with E units of it. Households each own an equal number of shares in the businesses and receives the profits they earn in the form of dividends.

There are two types of consumers: young and old. Let $\mathcal{I}_O \in [0,1]$ 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 the technology choices of businesses. We represent this difference as follows.

Assumption 1 (Preferences for technology)**.** *For old consumers, the consumption aggregate over varieties produced by businesses is:*

$$C(i) = \left(\int_0^J c(i,j)^{\rho} dj\right)^{\frac{1}{\rho}} \qquad \text{if } i \in \mathcal{I}_O. \tag{2}$$

Instead, for young consumers, the consumption aggregate over varieties is given by:

$$C(i) = \left(\int_0^J b(j)^{1-\rho} c(i,j)^{\rho} dj\right)^{\frac{1}{\rho}} \qquad \text{if } i \in \mathcal{I}_Y, \tag{3}$$

where $b(j) \geq 1$ depends on the technology adoption decision of business j.

Here, J is the number of varieties produced, which we index by j.[36] Moreover, $c(i,j)$ is the consumption of variety j by household i, and $\rho \in [0,1]$ determines the elasticity of substitution between varieties. The budget constraint of each household is:

$$\int_0^J p(j) c(i,j) dj + O(i) \leq E + \int_0^J \pi(j) dj. \tag{4}$$

where $p(j)$ is the price of variety j.[37] Maximization of (1) subject to (4) yields the usual demand curves:

$$\forall i \in \mathcal{I}_O, \quad c(i,j) \;=\; \left(\frac{p(j)}{P_o}\right)^{-\frac{1}{1-\rho}} C(i), \quad P_o \equiv \left(\int_0^J p(j)^{-\frac{\rho}{1-\rho}} dj\right)^{-\frac{1-\rho}{\rho}}, \tag{5}$$

[36] Each business produces a unique variety so we also use j to index businesses.

[37] We assume that businesses cannot price-discriminate between young and old.

$$\forall i \in \mathcal{I}_Y, \quad c(i,j) \;=\; b(j)\left(\frac{p(j)}{P_y}\right)^{-\frac{1}{1-\rho}} C(i), \quad P_y \equiv \left(\int_0^J b(j)p(j)^{-\frac{\rho}{1-\rho}}dj\right)^{-\frac{1-\rho}{\rho}}. \tag{6}$$

Because all young consumers and all old consumers make identical choices, we will omit the index i and refer instead to the consumption choices of old households by O_o, C_o and c_o, and to those of young households by O_y, C_y and c_y.

Businesses Each business is the monopolistic producer of its corresponding variety. Businesses all have the same unit cost of sales ξ and the same fixed operating cost ν. Finally, they face a cost of adopting technology level $\tilde{b}(j)$, which we model as follows.

Assumption 2 (Technology adoption costs)**.** *Choosing technology adoption level* $\tilde{b} \geq 1$ *requires* $\gamma(\tilde{b})$ *units of the numéraire good, where* $\gamma : [1, +\infty) \to \mathbb{R}^+$ *is a twice-differentiable, strictly increasing, and strictly convex function satisfying* $\gamma(1) = \gamma'(1) = 0$*.*

Because of externalities across businesses, the level of technology adoption $\tilde{b}(j)$ chosen by business j need not be the same as the effective impact of the technology on young households, $b(j)$. This is our third assumption.

Assumption 3 (Externalities in technology adoption)**.** *The technology adoption choice of business* j *affects young consumers' preferences as follows:*

$$\forall j \in [0, J], \quad b(j) = \bar{b}^{\theta}\, \tilde{b}(j), \quad \bar{b} \equiv \prod_{k \neq j} \tilde{b}(k)^{\frac{1}{J}}, \quad \theta \geq 0. \tag{7}$$

Profits for business j will therefore be given by:

$$\pi(j) = (p(j) - \xi)(\eta c_o(j) + (1-\eta)c_y(j)) - \gamma(\tilde{b}(j)) - \nu, \tag{8}$$

where $c_o(j)$ and $c_y(j)$ are given by the demand curves (5)-(6), and business j takes $\bar{b}$, the average level of technology adoption, as given.

Competitive equilibrium We will consider two types of competitive equilibria: those with a fixed number of businesses, J; and those with free-entry.

Definition 1 (Competitive equilibrium with fixed number of businesses)**.** *For a fixed number of businesses* J*, a competitive equilibrium is a set of prices and quantities such that (a) each household* i *maximizes utility (1) subject to their budget constraint (4); (b) each business* j *maximizes profits (8) subject to the demand curves (5)-(6) and taking* $\bar{b}$ *in Equation (7) as given.*

Definition 2 (Competitive equilibrium with free-entry)**.** *A competitive equilibrium with free-entry is a number of varieties* J *and a set of prices and quantities such that (a) each household* i *maximizes utility (1) subject to their budget constraint (4); (b) each business* j *maximizes profits (8) subject to (5)-(6) and (7); (c) the number of businesses* J *adjusts such that* $\inf_{j\in[0,J]} \pi(j) = 0$*.*

Appendix A.2 provides an analytical characterization of the two types of equilibria.

3.2 Discussion of key assumptions

Our model makes three non-standard assumptions, the point of each we now discuss.

Assumption 1: households' preferences for technology This assumption captures the idea that young consumers are more sensitive to businesses' technology offerings than old consumers. For the particular case of mobile payments, the focus of this paper, this assumption is consistent with the evidence presented in Section 2, which highlighted the quantitative importance of the negative relationship between age and the propensity to use the technology in India. As a consequence of this assumption, business j has the following market share of young consumers:

$$s_y(j) = b(j) \left(\frac{p(j)}{P_y} \right)^{-\frac{\rho}{1-\rho}}, \tag{9}$$

which has unit elasticity with respect to the technology choice $\tilde{b}(j) = \bar{b}^{-\theta} b(j)$ of business j. Thus business j technology adoption choice raises their market share of the young all else equal.[38]

Assumption 2: businesses' cost of technology adoption This assumption says that adopting the technology may create a burden on the business. One interpretation is that the technology may require workforce training to be deployed. Another interpretation is that businesses may be uncertain that the technology is reliable. This cost limits the scale of adoption.

Assumption 3: externalities in technology adoption We make this assumption because we use this model to study the adoption of a digital payments interface. Network externalities associated with digital payment systems have been recently documented in Parlour et al. (2022) (for banking payment infrastructure), Crouzet et al. (2023) (for retail payment interfaces), and Higgins (2024) (for debit cards and point of sales terminals). Our assumption is simple; in particular, externalities are one-sided in that young consumers' make no explicit technology adoption choice. The benefit of this simplicity is that it makes the model tractable while preserving the fundamental intuition that business decisions to adopt the payment technology, the object of our analysis, have positive spillovers across businesses that each business may fail to fully internalize.

Finally, we note two restrictions in the scope of our analysis. First, consumers' attitudes toward technology are treated as exogenous and determined by age. This restriction is significant. We make it deliberately so as to focus the analysis on how businesses adapt to consumer preferences (as implied by demographics), separate from how these preferences might evolve endogenously as a result to exposure to new technologies. However, we note that this assumption is consistent with the evidence, discussed in Section 2, that consumers preferences appear to be relatively rigid over the period we study, with limited evidence that adoption rates increase among the old over the period of time we study.

Second, the model has a single period. This rules out, in particular, businesses adapting to forecasted changes in demographic structure over time. This restriction could be relaxed, at the cost of more complicated exposition. However, we note that the free-entry equilibrium sheds light

[38]The choice of the Cobb-Douglas form $b(j)^{1-\rho}c(j)^{\rho}$ in the aggregator for young households is a normalization chosen to ensure that the market share $s_y(j)$ has unit elasticity with respect to $\tilde{b}(j)$.

on how the number of businesses might adjust over time in response to the introduction of the technology, given a particular demographic structure. The main benefit of these restriction is that we can compare the competitive equilibrium with efficient benchmarks and study optimal policy, which we do in Section 5.

3.3 Empirical predictions

The necessary first-order condition for profit maximization are:

$$p(j) = \frac{\xi}{\rho}, \tag{10}$$

$$(1-\rho)(1-\eta)P_y C_y \frac{\partial s_y(j)}{\partial \tilde{b}(j)} = \gamma'(\tilde{b}(j)). \tag{11}$$

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 $\frac{1}{\rho}$. The first-order condition for technology adoption equates its marginal benefit with its marginal cost. Aggregate spending by young households is $(1-\eta)P_y C_y$, and the business earns profits $(1-\rho)$ per dollar of sales to young consumers. Technology adoption serves to increase the business's market share of young consumers, all else equal, which is captured by the term $\partial s_y(j)/\partial \tilde{b}(j)$ on the right-hand side of Equation (11). Combining the two first-order conditions yields an alternative formulation for the optimal choice of technology:

$$(1-\eta)\frac{1-\rho}{\rho}\frac{c_y(j)}{\tilde{b}(j)} = \frac{\gamma'(\tilde{b}(j))}{\xi}. \tag{12}$$

The left-hand side can be interpreted as a private marginal rate of transformation between technology, $\tilde{b}(j)$, and quantity choice, $c_y(j)$, while the right-hand side is the relative marginal cost.

Because all businesses charge the same markup, the equilibrium is symmetric. As all businesses make identical choices, despite the fact that each of them attempts to attract more young customers by adopting the technology, their efforts cancel out. However, in this model technology adoption has general equilibrium effects. Appendix A.2 shows that, for any number of active business J, equilibrium household income is given by:

$$I = \frac{1}{1-(1-\rho)\alpha}\left(E - J\left(\gamma(\tilde{b}) + \nu\right)\right), \tag{13}$$

where $\tilde{b}$ is the equilibrium (private) technology adoption choice of each business. All else equal, more adoption lowers household income.

Prediction 1 (Demographics and technology adoption)**.** *In the competitive equilibrium, technology* *adoption $\tilde{b}$ is decreasing in the share of old consumers, η:*

$$\frac{\partial \tilde{b}}{\partial \eta} < 0.$$

This holds true both when the number of businesses is fixed or when there is free-entry, and regardless of whether there are network externalities (i.e. regardless of whether $\theta = 0$ or $\theta > 0$).

Proof. See Appendix A.2. ■

To see why this prediction holds, note that we can re-arrange condition (12) using the fact that demand from the young is $c_y = \alpha I/(pJ) = \alpha\rho I/(\xi J)$, to obtain:

$$\tilde{b}\gamma'(\tilde{b}) = (1-\eta)(1-\rho)\alpha\, \frac{I}{J}. \tag{14}$$

Differentiating both sides of Equation (14) with respect to η gives:

$$\frac{d}{d\tilde{b}}\left[\tilde{b}\gamma'(\tilde{b})\right]\frac{\partial \tilde{b}}{\partial \eta} = \underbrace{-(1-\rho)\alpha\frac{I}{J}}_{\text{effect on size of young consumer base}} + \underbrace{(1-\eta)(1-\rho)\frac{\alpha}{J}\frac{\partial I}{\partial \eta}}_{\text{income effect}}$$

The function on the left-hand side is strictly increasing, because of Assumption 1. On the right-hand side, the first term is a direct effect of the change in demographics on the marginal benefit from technology adoption for each business. The second term is an indirect general equilibrium effect — all else equal, a demographic shift changes technology adoption choices, which in turns affect corporate profits and household income. Using Equation (13) and re-arranging, we obtain:

$$\overbrace{\left\{\frac{d}{d\tilde{b}}\left[\tilde{b}\gamma'(\tilde{b})\right] + (1-\eta)\frac{(1-\rho)\alpha}{1-(1-\rho)\alpha}\gamma'(\tilde{b})\right\}}^{>0}\frac{\partial \tilde{b}}{\partial \eta} = \underbrace{-(1-\rho)\frac{\alpha I}{J}}_{<0},$$

establishing the result.

Prediction 2 (Adoption costs)**.** *Suppose the adoption cost function can be written as:*

$$\gamma(\tilde{b}) = \frac{1}{2}\omega(\tilde{b}-1)^2. \tag{15}$$

An increase in ω leads to a lower rate of technology adoption among businesses; and the effect is weaker, the higher the share of old consumers, η:

$$\frac{\partial \tilde{b}}{\partial \omega} < 0, \quad \frac{\partial^2 \tilde{b}}{\partial \omega \partial \eta} > 0.$$

Proof. See Appendix A.2. ■

Thus the model implies that, when the share of old consumers is higher (that is, when η increases), there is a weaker incentive for individual businesses to invest in technology, in order to gain market share.[39] Moreover, higher technology adoption costs (a higher ω) weaken this incentive. In the following section, we will test these predictions by contrasting the technology adoption choices of businesses that face consumer bases with different demographics.

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.[40] Our analysis will focus on examining how the adoption of the fintech company's terminals by stores has evolved over time.[41]

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 allow individuals to directly use mobile payment options, for instance paying using UPI through any supporting apps.[42] Lastly, although our fintech company is sizable, it represents just one

[39]Note 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.

[40]In 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.

[41]We 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.

[42]In 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

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.[43]

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 a 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).[44] Finally, we manually collected a list of universities in the country as of 2019, and mapped each university to its official pincode.[45] 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 availability 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

most of QR transactions are UPI. Appendix A.1.2 contains more details on the data provided by our fintech partner and on the different POS offerings.

[43]Therefore, 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.

[44]A 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.

[45]Data Appendix A.1.3 discusses the construction of this data.

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.[46]

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:

$$y_{dt} = \alpha_d + \alpha_t + \beta \ (\text{AgeStructure}_d \times 1_{\{t \geq t_0\}}) + \Gamma_t^{'} \mathbf{X}_d + \epsilon_{dt}, \tag{16}$$

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.[47] When plotting the dynamic effects, we normalize the last month of the pre-period (i.e., April 2019) to zero in the following specification:

$$y_{dt} = \alpha_d + \alpha_t + \sum_{k=-6, k \neq -1}^{k=+6} \beta_k \ (\text{AgeStructure}_d \times 1_{\{t = t_0 + k\}}) + \Gamma_t^{'} \mathbf{X}_d + \epsilon_{dt}. \tag{17}$$

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

[46] If 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.

[47] To 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 the district from the Census.[48] 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 16. 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.[49]

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 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

[48] To 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.

[49] The same result is also presented in Table 2, column 3.

transactions on their platform.[50]

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.[51] 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.[52] In Table A-1, 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

[50] In 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.

[51] Variables 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.

[52] As 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.

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.[53] 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.[54] 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 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

[53] We 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.

[54] Our 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).

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 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.[55] 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.

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.[56] With this data, we are able to identify the list

[55] For 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.

[56] Data Appendix A.1 outlines the methodology employed to gather this information. It's important to acknowledge

of pincodes in India that host at least one university. Then, we estimate a differences-in-differences estimator of the following form:

$$y_{pt} = \alpha_{dt} + \alpha_p + \sum_{k=-6, k\neq -1}^{k=+6} \gamma_k \ \left(1\{Univ = 1\}_p \times 1_{\{t=t_0+k\}}\right) + \nu_{pt} \tag{18}$$

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.[57] 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 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

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.

[57] One 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.

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.[58] 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 provide 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).[59]

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.[60]

5 Subsidizing technology adoption with heterogeneous consumers

Having validated empirically the model's basic predictions, we now compare its competitive equilibrium (CE) to the first best (FB) allocation to understand when and how policy should encourage technology adoption when businesses face consumers with heterogeneous valuations of the technol-

[58] To 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.

[59] In 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).

[60] Our 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.

ogy, as is the case in the Indian payments context.[61] Two forces potentially depress adoption in CE, both of which could interact with heterogeneity in consumer valuations: (i) monopoly pricing (constant markup), which restricts output and thus weakens incentives to adopt a complementary technology; and (ii) network externalities across businesses, which are not internalized. We first analyze the fixed-J ("short run") case, then the free-entry ("medium run") case. Throughout, we use simple closed forms to show what the planner does with a single instrument—a constant marginal subsidy to adoption expenses—and why it restores adoption but not output in the short run, and raises adoption while lowering entry in the long run. Proofs are in Appendix A.3.

5.1 The short run: fixed number of businesses

1. First-best benchmark

Definition 3. *The first-best (FB) allocation are values for* $\left\{O_o, O_y, c_o, c_y, \tilde{b}\right\}$ *that maximize the welfare criterion:*

$$W = \eta W_o + (1-\eta) W_y = \eta \log\left(O_o^{1-\alpha} C_o^{\alpha}\right) + (1-\eta) \log\left(O_o^{1-\alpha} C_o^{\alpha}\right). \tag{19}$$

subject to the resource constraint:

$$\eta O_o + (1-\eta) O_y + \xi\left(\eta c_o + (1-\eta) c_y\right) + J\left(\gamma(\tilde{b}) + \nu\right) \leq E. \tag{20}$$

where $C_o = J^{\frac{1}{\rho}} c_o$, $C_y = J^{\frac{1}{\rho}} \tilde{b}^{(1+\theta)\frac{1-\rho}{\rho}} c_y$.

Result 1 (First-best allocation)**.** *In the first-best allocation, the planner chooses* $(c_y, \tilde{b})$ *so that:*

$$(1+\theta)(1-\eta)\frac{1-\rho}{\rho}\frac{c_y}{\tilde{b}} = \frac{\gamma'(\tilde{b})}{\xi}. \tag{21}$$

Moreover, the first-best allocation is consistent with utility maximization of consumers if prices are equal to the marginal cost of production, $p = \xi$*, so that business profits are strictly negative. Finally, in the first-best allocation, technology adoption is higher than in the competitive equilibrium (CE), even without externalities (*$\theta = 0$*):*

$$\tilde{b}_{FB} > \tilde{b}_{CE}. \tag{22}$$

To understand this result, start by recalling that in the (symmetric) competitive equilibrium,

[61] In Appendix A.3, we also compare the CE allocation to another benchmark, the "constrained optimal" (CO) allocation (Dixit and Stiglitz 1977; Dhingra and Morrow 2019). This allocation maximizes welfare subject to businesses breaking even. The comparison between the CO and the CE is largely similar to our comparison of the FB and the CE, with one notable exception: under free-entry, in the absence of externalities ($\theta = 0$), the CE and the CO exactly coincide, including with respect to the number of businesses, J. This is a version of the classic Dixit and Stiglitz (1977) constrained efficiency result applied to our framework, which includes heterogeneous consumers and a technology choice. Our focus is on optimal adoption subsidies. With adoption subsidies, the set of feasible allocations in the CE contains some that would violate the positive profit condition without subsidies, and are therefore not feasible in the CO. Because our subsidies are financed from lump-sum taxation of the household, they are within the feasible set of the FB planning problem. This motivates our to use the FB as our benchmark.

the first-order condition relating output and technology adoption is:

$$(1-\eta)\frac{1-\rho}{\rho}\frac{c_y}{\tilde{b}} = \frac{\gamma'(\tilde{b})}{\xi}, \tag{23}$$

Suppose first that there are no spillovers in technology adoption across businesses ($\theta = 0$). Then Equation (21) and (23) are the same. The social and private marginal rate of substitution between output and technology are identical, so that for a given level of output c_y, the corresponding adoption choice in the CE would be first-best efficient.

Yet Result 1 states that even in this case, technology adoption is too low relative to first-best. This is because of the other source of inefficiency, monopolistic competition. Businesses in the CE restrict quantity so as to be able to charge higher prices. Figure 6, right panel, reports the level of production in the CE relative to the first-best when $\theta = 0$, and shows that the ratio is always below 1. Since output and technology are complements, the quantity restriction leads to under-investment in technology. Thus monopoly distortions affect the level of technology adoption, giving a potential role for subsidies to restore efficiency even in the absence of network externalities.

When $\theta > 0$, *both* the monopoly distortion and the un-internalized network spillovers across businesses depress the competitive level of technology adoption relative to the first-best. This is shown in the left panel of Figure 6, which reports the ratio $\tilde{b}_{CE}/\tilde{b}_{FB}$. The figure also shows that this wedge is increasing with respect to θ, for any given demographic composition η, highlighting that the un-internalized network spillovers amplify the under-investment problem.

2. Optimal adoption subsidy

Consider encouraging the adoption of the technology by using a constant marginal subsidy rate to the adoption cost, $\gamma(\tilde{b})$.[62] Firm profits become:

$$\pi = (p-\xi)(\eta c_o + (1-\eta)c_y) - ((1-\tau)\gamma(\tilde{b}) + \nu). \tag{24}$$

Moreover, assume that the subsidy is financed via lump-sum taxation so that income become:

$$I = E + J\left(\pi - \tau\gamma(\tilde{b})\right). \tag{25}$$

Appendix A.2 characterizes the competitive equilibrium with an adoption subsidy. We ask the following question: what value of the subsidy maximizes the welfare criterion (19)?

Result 2 (Optimal adoption subsidy)**.** *The optimal adoption subsidy is given by:*

$$\tau^* = \begin{cases} 0 & \text{if} \quad \eta = 1 \\ 1 - \dfrac{1}{1+\theta}\dfrac{\rho}{1-(1-\rho)\alpha} & \text{if} \quad 0 \le \eta < 1 \end{cases} \tag{26}$$

[62] In a model with business taxes, this could capture a rate of deductibility of expenses related to technology adoption from taxable business income.

In particular, (a) τ^* is strictly positive even in the absence of externalities; (b) it increases with the strength of externalities, θ; (c) it is independent of the share of old consumers, η; (d) it is independent of the features of technology adoption costs, $\gamma(.)$. Under the optimal adoption subsidy, adoption reaches its first-best level, but output does not:

$$\tilde{b}_{CE}(\tau^*) = \tilde{b}_{FB}, \quad c_{y,CE}(\tau^*) < c_{y,FB}. \tag{27}$$

This result has several interesting features. First, the optimal subsidy is positive even without externalities ($\theta = 0$). This is because monopoly distortions alone lead to under-adoption, but highlights how externalities are not necessary to justify technology adoption subsidies.

Second, the optimal adoption subsidy is independent of the share of young consumers, so long as that share is non-zero. To understand why, consider the two first-order conditions characterizing the optimal ratio of technology adoption to output in the CE with a subsidy and in the FB:

$$\begin{aligned} (1-\eta)\frac{1-\rho}{\rho}\frac{c_{y,CE}}{\tilde{b}_{CE}} &= (1-\tau)\gamma'(\tilde{b}_{CE}) \\ (1+\theta)(1-\eta)\frac{1-\rho}{\rho}\frac{c_{y,FB}}{\tilde{b}_{FB}} &= \gamma'(\tilde{b}_{FB}) \end{aligned}$$

Both the planner and the market weigh the benefits of the technology by the fraction of young consumers, $1-\eta$; and both agree on the marginal cost of adoption. Suppose that there were a subsidy level that equalized adoption between the CE and FB; then taking ratios of the first-order conditions above, one obtains:

$$\frac{c_{y,CE}}{c_{y,FB}} = (1-\tau)(1+\theta). \tag{28}$$

Furthermore, because markups are constant, both income and prices in the CE are simple multiplicative function of those corresponding to the FB.[63] As a result, under identical technology adoption choices:

$$\frac{c_{y,CE}}{c_{y,FB}} = \frac{\rho}{1-(1-\rho)\alpha}, \tag{29}$$

which is independent of demographics or of the shape of the technology adoption costs. Comparing (28) and (29), we see that the subsidy consistent with raising the CE adoption level to its FB counterpart must be independent of demographics or the shape of technology adoption costs.

The reason why the planner setting the subsidy uses it to restore adoption to its first-best level is that the subsidy itself does not affect the markup chosen by businesses, which remains equal to $1/\rho$ even with the subsidy. Thus the planner cannot target directly the monopoly distortion, but it can mitigate its effects on technology adoption. This is also the reason why the subsidy fails to restore the first-best allocation and consumption remains distorted downward.[64]

[63] This property would hold even without technology choice in the model, and only depends on CES preferences.

[64] Adding a proportional production subsidy, or a tax on sales, as an additional instrument, would allow the planner to target both the adoption level and the output level separately and restore efficiency.

A lower bound on the optimal adoption subsidy is given by:

$$\tau^* \geq \underline{\tau}^* = 1 - \frac{\rho}{1-(1-\rho)\alpha}. \tag{30}$$

Note that this lower bound can be written as:

$$\underline{\tau}^* = \frac{(\frac{1}{\rho}-1)(1-\alpha)}{1+(\frac{1}{\rho}-1)(1-\alpha)}. \tag{31}$$

Under the model described above, given data on total variable costs $\mathcal{C} = J\xi c_y$, total business sales $\mathcal{S} = Jpc_y$, and total household expenditures $\mathcal{X} = \mathcal{S} + (1-\eta)O_y + \eta O_o$, this lower bound can be computed using the fact that $\alpha = \frac{\mathcal{S}}{\mathcal{X}}$ and $\rho = \frac{\mathcal{C}}{\mathcal{S}}$, yielding:

$$\underline{\tau}^* = \frac{(\mathcal{S}-\mathcal{C})(\mathcal{X}-\mathcal{S})}{\mathcal{C}\mathcal{X}+(\mathcal{S}-\mathcal{C})(\mathcal{X}-\mathcal{S})}. \tag{32}$$

This lower bound (which coincides with the optimal subsidy when $\theta = 0$) can be viewed as the Pigouvian subsidy that exactly offsets the effects of the monopoly distortion on adoption; network externalities then further add subsidy incentives beyond that lower bound.[65]

5.2 The medium-run: free entry

With endogenous entry, an additional inefficiency arises in the competitive equilibrium: businesses cannot perfectly internalize the consumer surplus associated with the greater variety (the "love-for-variety" effect) that their entry creates. As we will show, this demand externality can also interact with adoption decisions.

1. First-best benchmark

Definition 4. *The first-best (FB) allocation are quantities* $\left\{O_o, O_y, c_o, c_y, \tilde{b}\right\}$ *and a number of businesses* J *that maximize the welfare criterion (19) subject to the resource constraint (20).*

Result 3 (First-best allocation)**.** *In the first-best allocation, the planner chooses an adoption rate that satisfies:*

$$(1-\eta)(1+\theta)(\gamma(\tilde{b})+\nu) = \gamma'(\tilde{b})\tilde{b}. \tag{33}$$

If there are no externalities ($\theta = 0$), technology adoption and consumption are equal to their competitive equilibrium level, though entry is higher in the first-best than in the competitive equilibrium.

[65] Letting $L = 1-\rho$, on can write $\underline{\tau}^* = \frac{L(1-\alpha)}{1-\alpha L}$. The Lerner index L captures the "tax" that markups place on the output of businesses on the model, and the term $(1-\alpha)L$ reflects the ability of consumers to substitute into the outside good as a result of that tax. The denominator $1-\alpha L$ reflects higher income from profits rebated to households, which partially offsets the "tax" in the numerator. The resulting optimal subsidy exactly offsets the effect of this mix of two distortions coming from monopoly pricing on adoption.

With externalities ($\theta > 0$), the first-best allocation features more technology adoption and production than the competitive equilibrium:

$$\tilde{b}_{FB} > \tilde{b}_{CE}, \quad c_{y,FB} > c_{y,CE}. \tag{34}$$

The first-best allocation can also feature less entry than the competitive equilibrium for a sufficiently large value of the externality parameter θ.

To understand this result, note that the competitive equilibrium level of adoption with free-entry is given by:

$$(1-\eta)(\gamma(\tilde{b}) + \nu) = \gamma'(\tilde{b})\tilde{b}. \tag{35}$$

Equation (35) can be thought of as follows. With free-entry, the markup revenue from production of each variety must be equal to the total overhead costs $\gamma(\tilde{b}) + \nu$. As explained in Section 3, the elasticity of markup revenue to adoption is constant equal to $1 - \eta$, so that the equilibrium marginal benefit of technology adoption will equal $(1 - \eta)(\gamma(\tilde{b}) + \nu)/\tilde{b}$. Thus the equation above equates the marginal benefits from technology adoption to their marginal cost. In the absence of network externalities ($\theta = 0$), this equation coincides with the planner's first-order condition, (33), leading to an efficient outcome for adoption. Because the planner and the market also have the same marginal rate of substitution between output and technology when $\theta = 0$, an efficient level of adoption also coincides with an efficient level of production.

Note, however, that even without externalities ($\theta = 0$), the competitive equilibrium allocation does not lead to the first-best level of entry. This is because of the love-for-variety effect: the planner would value more consumer surplus from additional varieties, but businesses cannot capture that surplus once the break-even condition binds.[66] Figure 7, right panel, illustrates this — when $\theta = 0$, entry remains inefficiently low.

Finally, Result 3 states that entry might be inefficiently *high* in the CE when adoption externalities are strong enough. This is shown, numerically, in the right panel of Figure 7. With positive network externalities, the first-best planner would want to adopt the technology more broadly. This increases the overhead resource costs $\gamma(\tilde{b}) + \nu$ associated with the creation of each variety. Despite the love-for-variety effect, the first-best planner is willing to tolerate somewhat lower entry in order to reduce these welfare costs, effectively substituting variety for broader technology adoption.

2. Optimal adoption subsidy

Let $W(\tau; \theta, \eta)$ be the welfare criterion from Equation (19) under the allocation corresponding to the competitive equilibrium with free entry and with an adoption subsidy equal to $\tau \geq 0$. We denote by $\tau^{**}(\eta, \theta)$ the subsidy that maximizes $W(\tau; \theta, \eta)$.

[66] As mentioned above, we show in Appendix A.3 that the CE level of entry is constrained-efficient, that is, coincides with the level of entry chosen by a planner bound to meet a zero-profit condition.

Result 4 (Optimal adoption subsidy with free-entry)**.** *For all $\eta < 1$, a small adoption subsidy improves welfare:*

$$\frac{\partial W}{\partial \tau}(0; \eta, \theta) > 0 \quad \forall \eta \in [0, 1), \quad \forall \theta \geq 0, \tag{36}$$

even in the absence of externalities. The optimal adoption subsidy in the competitive equilibrium with free-entry is therefore strictly positive: $\tau^{**}(\eta, \theta) > 0$.

Moreover, absent externalities, the marginal welfare gain from a small adoption subsidy is smaller, the larger the share of old consumers:

$$\frac{\partial}{\partial \eta}\left(\frac{\partial W}{\partial \tau}(0; \eta, \theta)\right) < 0 \qquad \textit{when } \theta = 0. \tag{37}$$

To understand this result, the key is to note that With free entry, an adoption subsidy is a second-best tool: it raises adoption b, but also increases overhead costs $\gamma(\tilde{b}) + \nu$ for each business, thus potentially lower entry J, even when externalities are absent ($\theta = 0$).

First consider the case $\theta = 0$. An adoption subsidy is always desirable, *even though* adoption in the CE without subsidy is already at its first-best level, as indicated by Result 3. Since adoption increases with the subsidy, it will therefore be *above* its first-best level. This is illustrated in the bottom middle panel of Figure 7, which reports the level of adoption under the optimal subsidy $\tau = \tau^{**}$ relative to the CE without subsidy.[67] The bottom right panel of Figure 7 also reports the entry rate under the optimal subsidy, and shows that it is *lower than* the entry rate in the CE without subsidy, and therefore also lower than the entry rate in the first-best.

Lower entry under the optimal subsidy might be surprising. In the CE, entrants cannot appropriate the surplus from more varieties once the zero-profit condition binds; one might have expected policy to aim at relaxing that constraint. However, with an adoption subsidy the planner cannot directly increase the number of varieties. Instead, an adoption subsidy raises overhead costs $\gamma(b) + \nu$, and tightens free entry, lowering the number of businesses. At the optimum, the gains from higher adoption outweigh the loss from fewer varieties. This is a second-best outcome: with an entry subsidy, the planner could instead target entry directly, and leave adoption unchanged.

Result 3 also shows that, even when $\theta = 0$, a small subsidy has a larger effect on welfare when the share of old consumers is smaller. A direct implication of this second point is that even with free-entry, a planner should subsidize adoption everywhere, but more aggressively in places where consumers are younger. This stands in contrast with the optimal "short-run" subsidy, which is entirely independent of demographics. Underscoring this point, the bottom left panel of Figure 7 shows that the optimal subsidy $\tau^{**}(\eta)$ is decreasing with respect to η even without externalities. As discussed above, the effect of the adoption subsidy is primarily to increase adoption (the welfare effects of which are limited to the young), at the expense of reducing variety (the welfare effects of which are shared by the young and the old).

In the case $\theta > 0$, the optimal subsidy is also positive. But in that case, adoption in the CE

[67] The optimal subisdy rate τ^{**} does not admit an analytical characterization, so we use the numerical example in Figure 7 to discuss the properties of the CE under the optimal subsidy rate.

without subsidy is too low relative to first-best, as businesses cannot fully internalize the consumer surplus that network externalities create. The resulting optimal subsidy is naturally higher, as shown in the bottom left panel of Figure 7. Because the social returns to higher adoption are larger, the planner is also willing to tolerate a larger reduction in entry relative to the CE; that is, with only an adoption subsidy, the planner substitutes lower variety with a higher rate of technology adoption. That effect is stronger, the larger the share of young consumers, as indicated by the bottom right panel of Figure (7).

In summary, with free-entry, subsidizing adoption is always optimal, but comes at the expense of reduced entry relative to the competitive equilibrium level. Moreover, because the benefits of higher adoption entirely fall on the young, the subsidy is increasing in the share of young consumers.

6 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. 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. Finally, we showed that, with heterogeneous consumer valuations of the technology, adoption by businesses is generically too low compared to first-best, though the resulting optimal adoption subsidy need not depend on the details of technology adoption costs or the demographic structure of the consumer base.

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.[68]

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

[68]We 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.

countries.[69] 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. In our context, adoption is significantly lower than optimal, and policymakers can potentially intervene to reduce this problem. However, the size of the subsidy is independent of the composition of consumers and adoption costs.

[69] 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. Duflo, 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.

Dhingra, S. and J. Morrow (2019). Monopolistic competition and optimum product diversity under firm heterogeneity. *Journal of Political Economy 127*(1), 196–232.

Dixit, A. K. and J. E. Stiglitz (1977). Monopolistic competition and optimum product diversity. *The American economic review 67*(3), 297–308.

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.

Higgins, S. (2024). Financial technology adoption: Network externalities of cashless payments in mexico. *American Economic Review 114*(11), 3469–3512.

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.

Parlour, C. A., U. Rajan, and J. Walden (2022). Payment system externalities. *The Journal of Finance 77*(2), 1019–1053.

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

(a) Volume of transactions

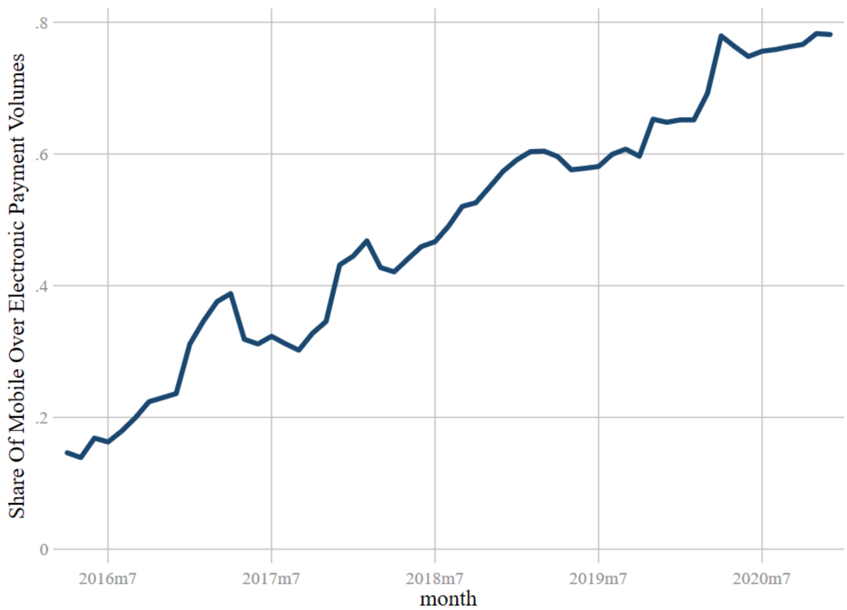

(b) Value of transactions

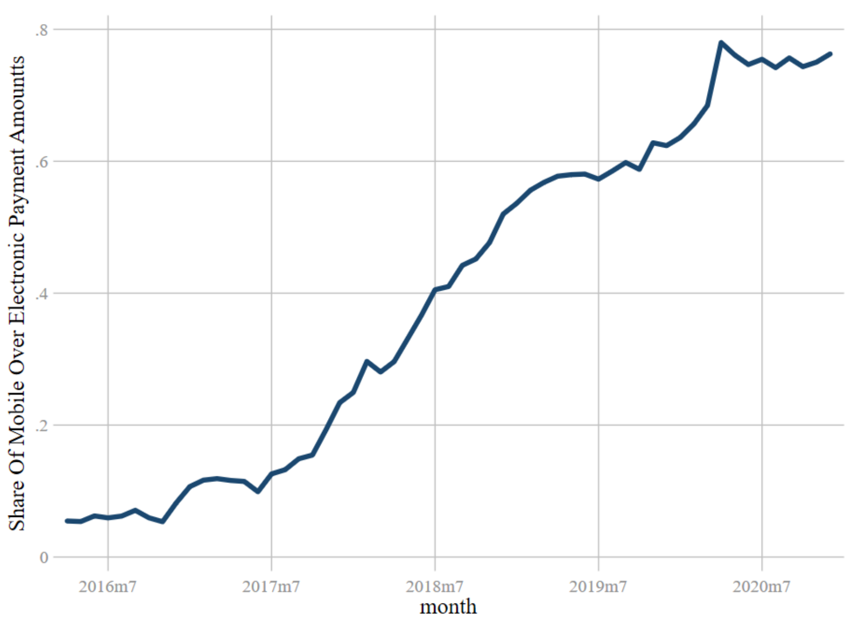

Notes: These two figures plot the share of electronic transactions that are done using mobile (at monthly level). Panel (a) reports the measure based on volume of transactions, while panel (b) examines the value of transaction. Electronic transactions are defined as the sum of UPI, mobile wallets, and 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

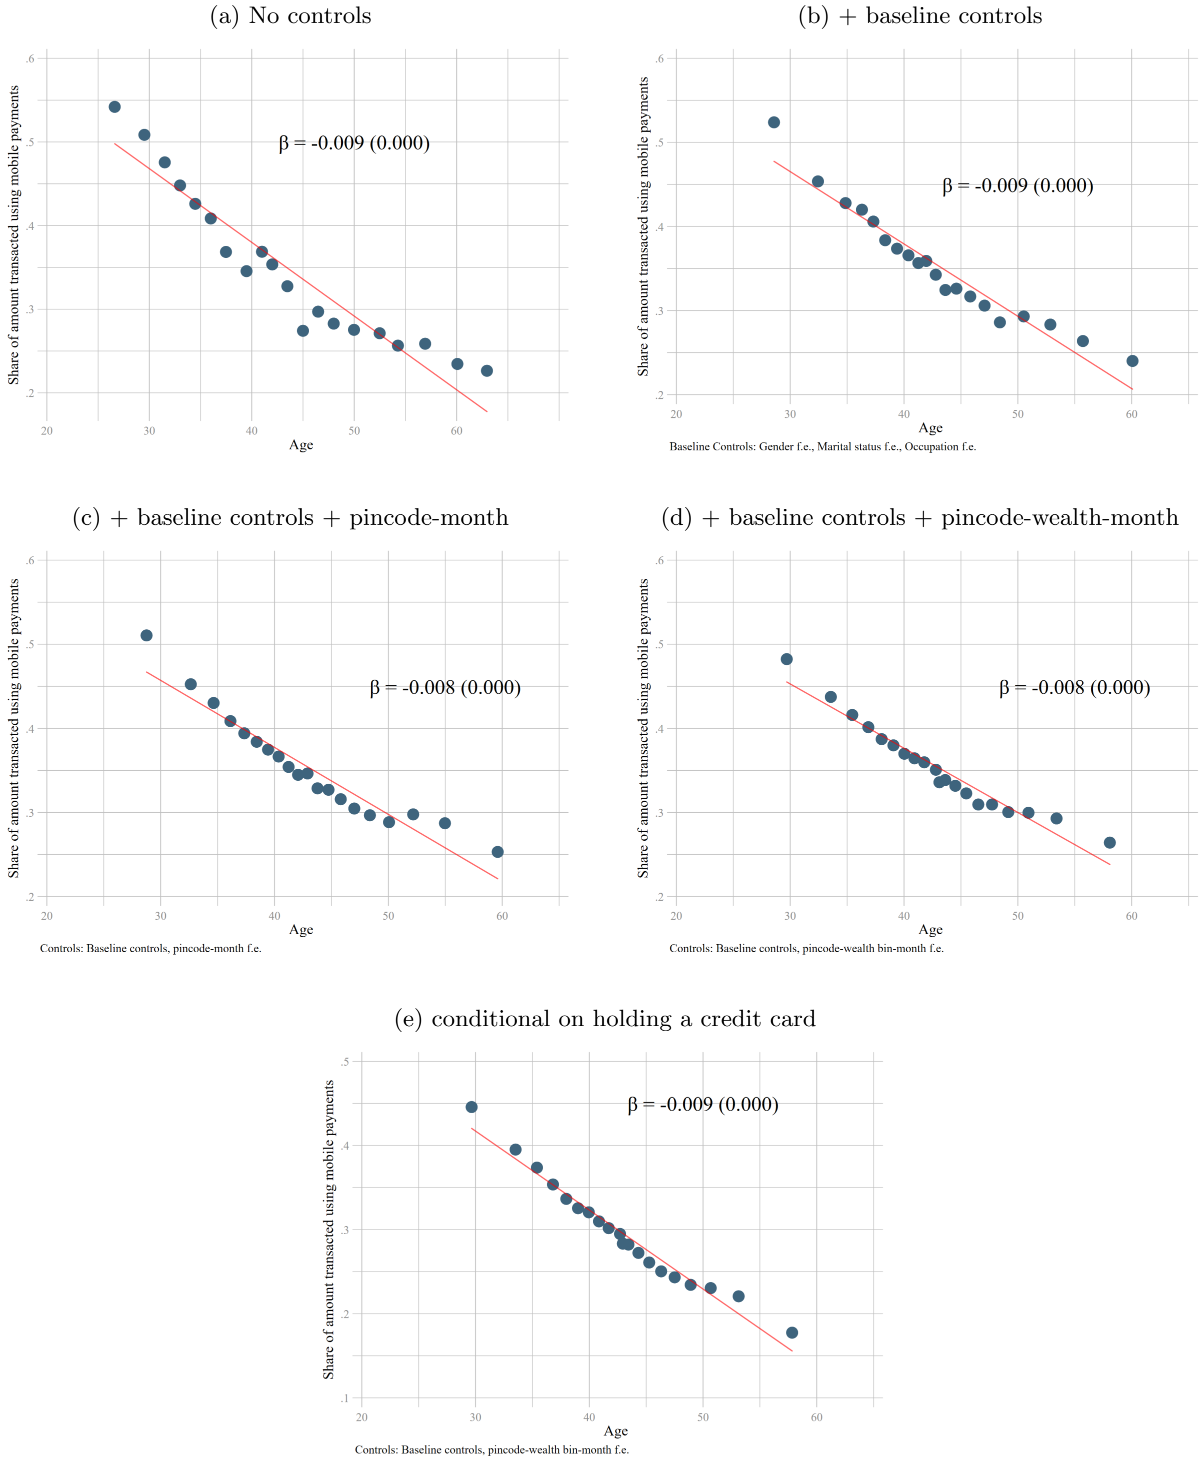

Notes: The figure plots the share of the amount transacted by households using mobile payments in the bank-level transaction data. All figures report the average of the share across 20 age bins and the slope of the line is reported in each panel. Panel (a) reports the mean with no controls. Panel (b) 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))

(a) baseline specification

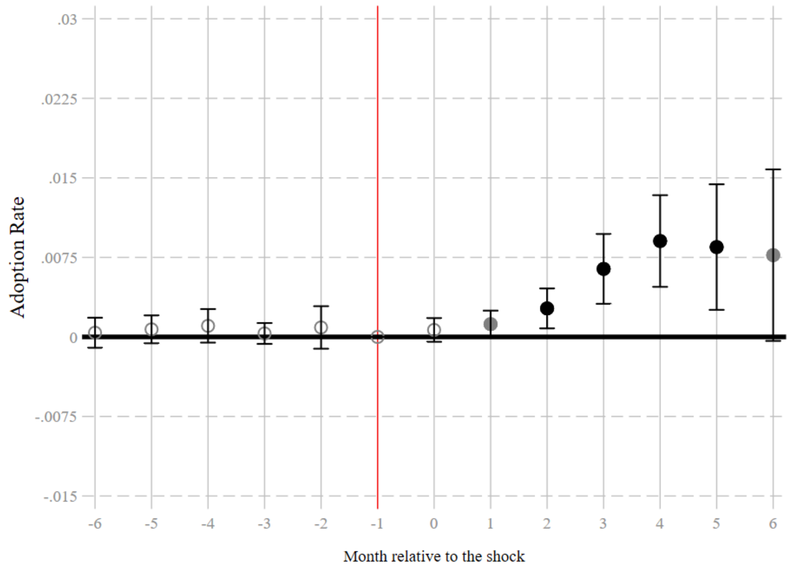

(b) all controls

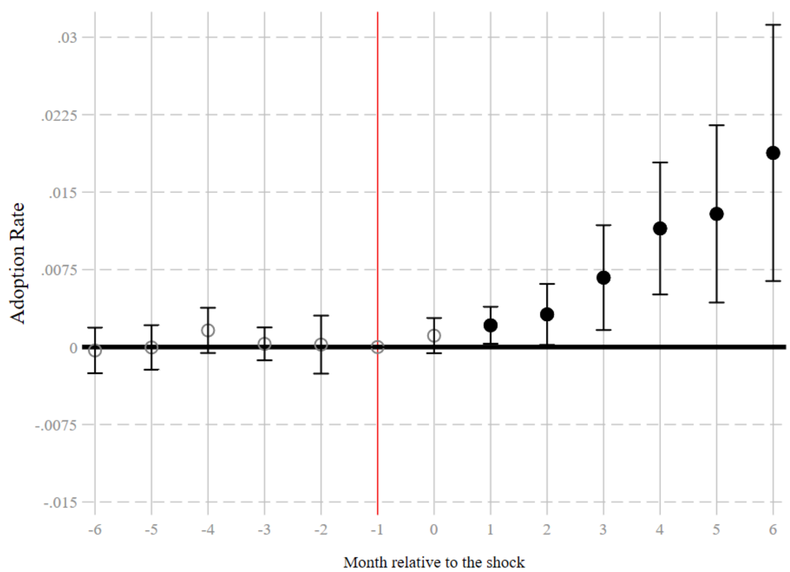

Notes: The figure plots the dynamic treatment effects of age structure on adoption. The dependent variable is the number of stores that adopted our fintech company in month t and district d, scaled by the total number of firms in the districts (in hundreds) measured by the Census. The period considered 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 17. 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

(a) baseline specification

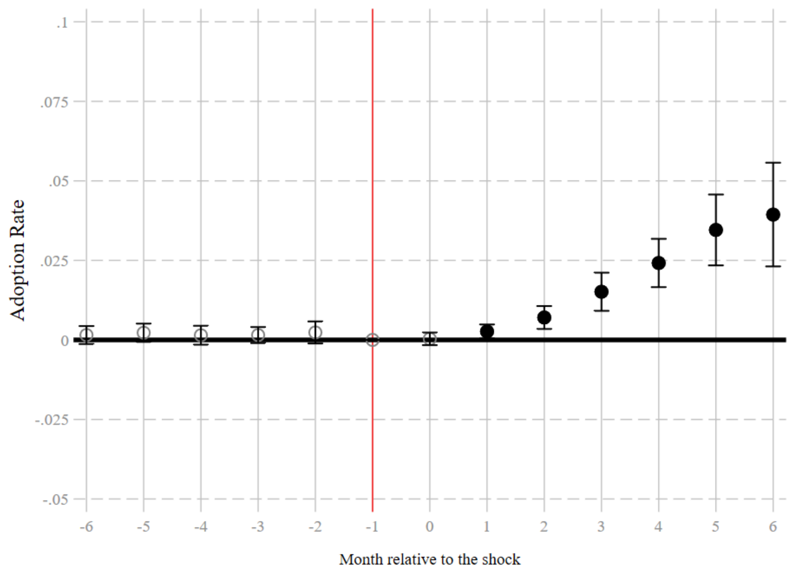

(b) all controls

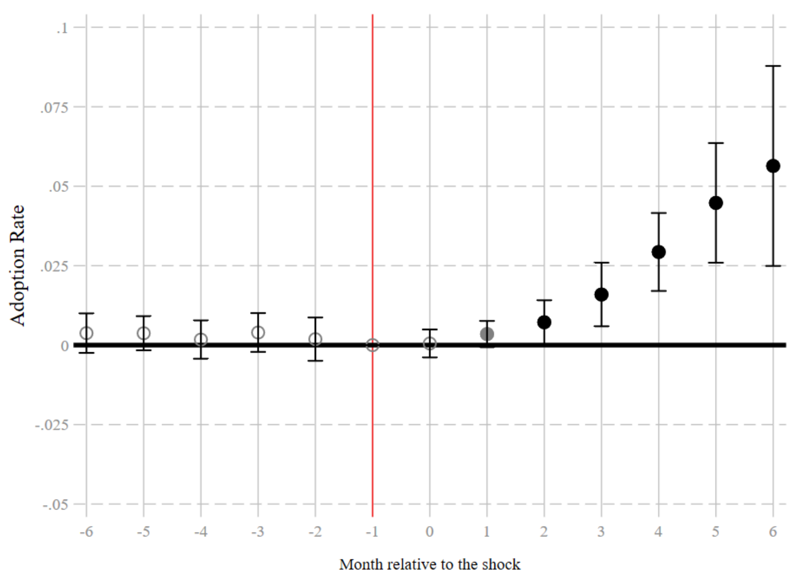

Notes: The figure plots the dynamic treatment effects of age structure on adoption, where the share of young adults is instrumented by a function of the historical sex-ratio, as described in the paper. The dependent variable is the number of stores that adopted our fintech company in month t and district 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 17. 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

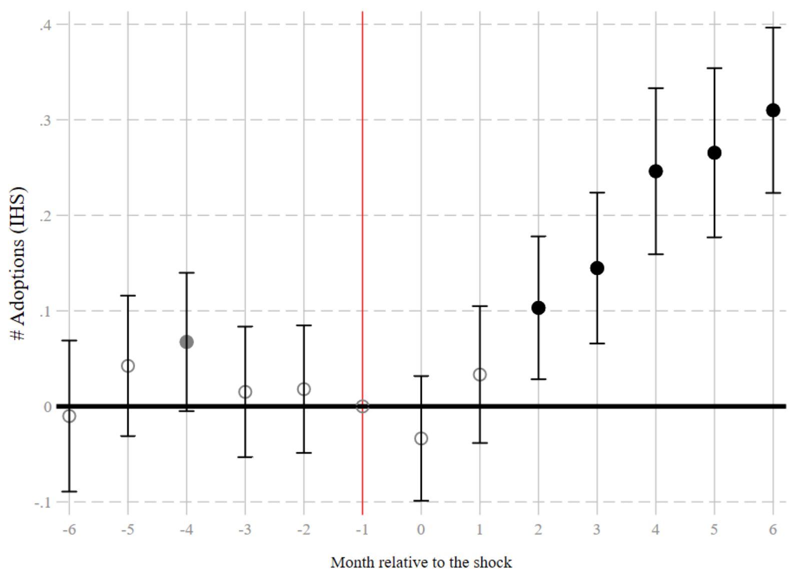

Notes: The figure plots the dynamic treatment effects of the presence of a university on the adoption of our fintech company. The dependent variable is the (inverse hyperbolic sine transformation of) the number of stores that adopted our fintech company at pincode-level in a month. The graphs report the coefficients γ_k from specification 18, 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 6: Competitive equilibrium (CE) and first-best (FB) with fixed number of businesses.

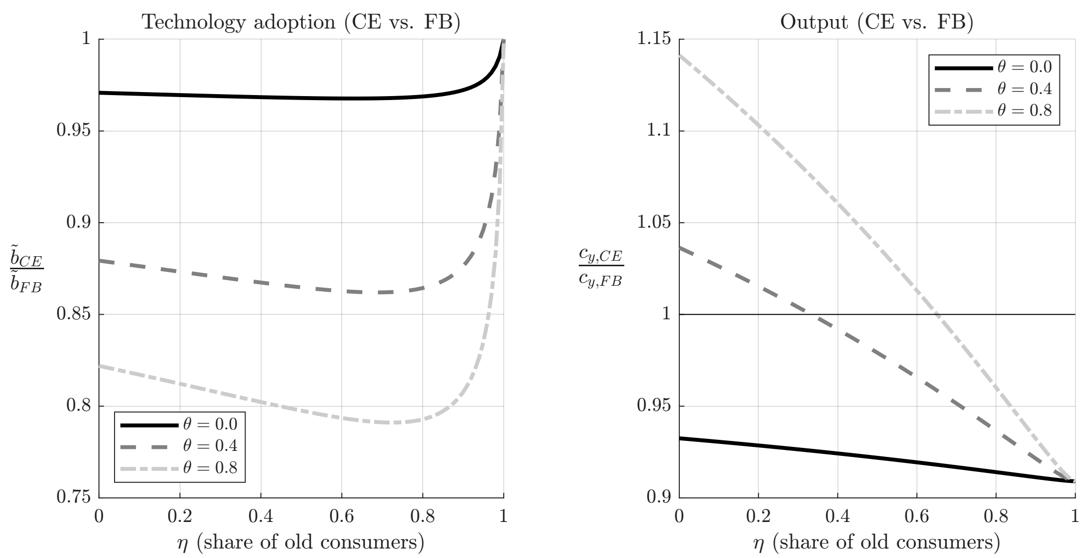

Notes: The left panel reports the ratio of adoption in the CE relative to FB. The right panel reports the ratio of output in the CE with $\tau = 0$ (no subsidy) relative to the FB. The horizontal axis is the fraction of old consumers. The calibration used is $E = 1$, $J = 2$, $\alpha = 0.9$, $\rho = 0.5$, $\nu = 0.1$, $\xi = 0.1$, and $\gamma(b) = \frac{1}{2}\omega(b-1)^2$ with $\omega = 0.025$. Results are reported for different degrees of network externalities.

Figure 7: Competitive equilibrium (CE) and first-best (FB) with endogenous entry.

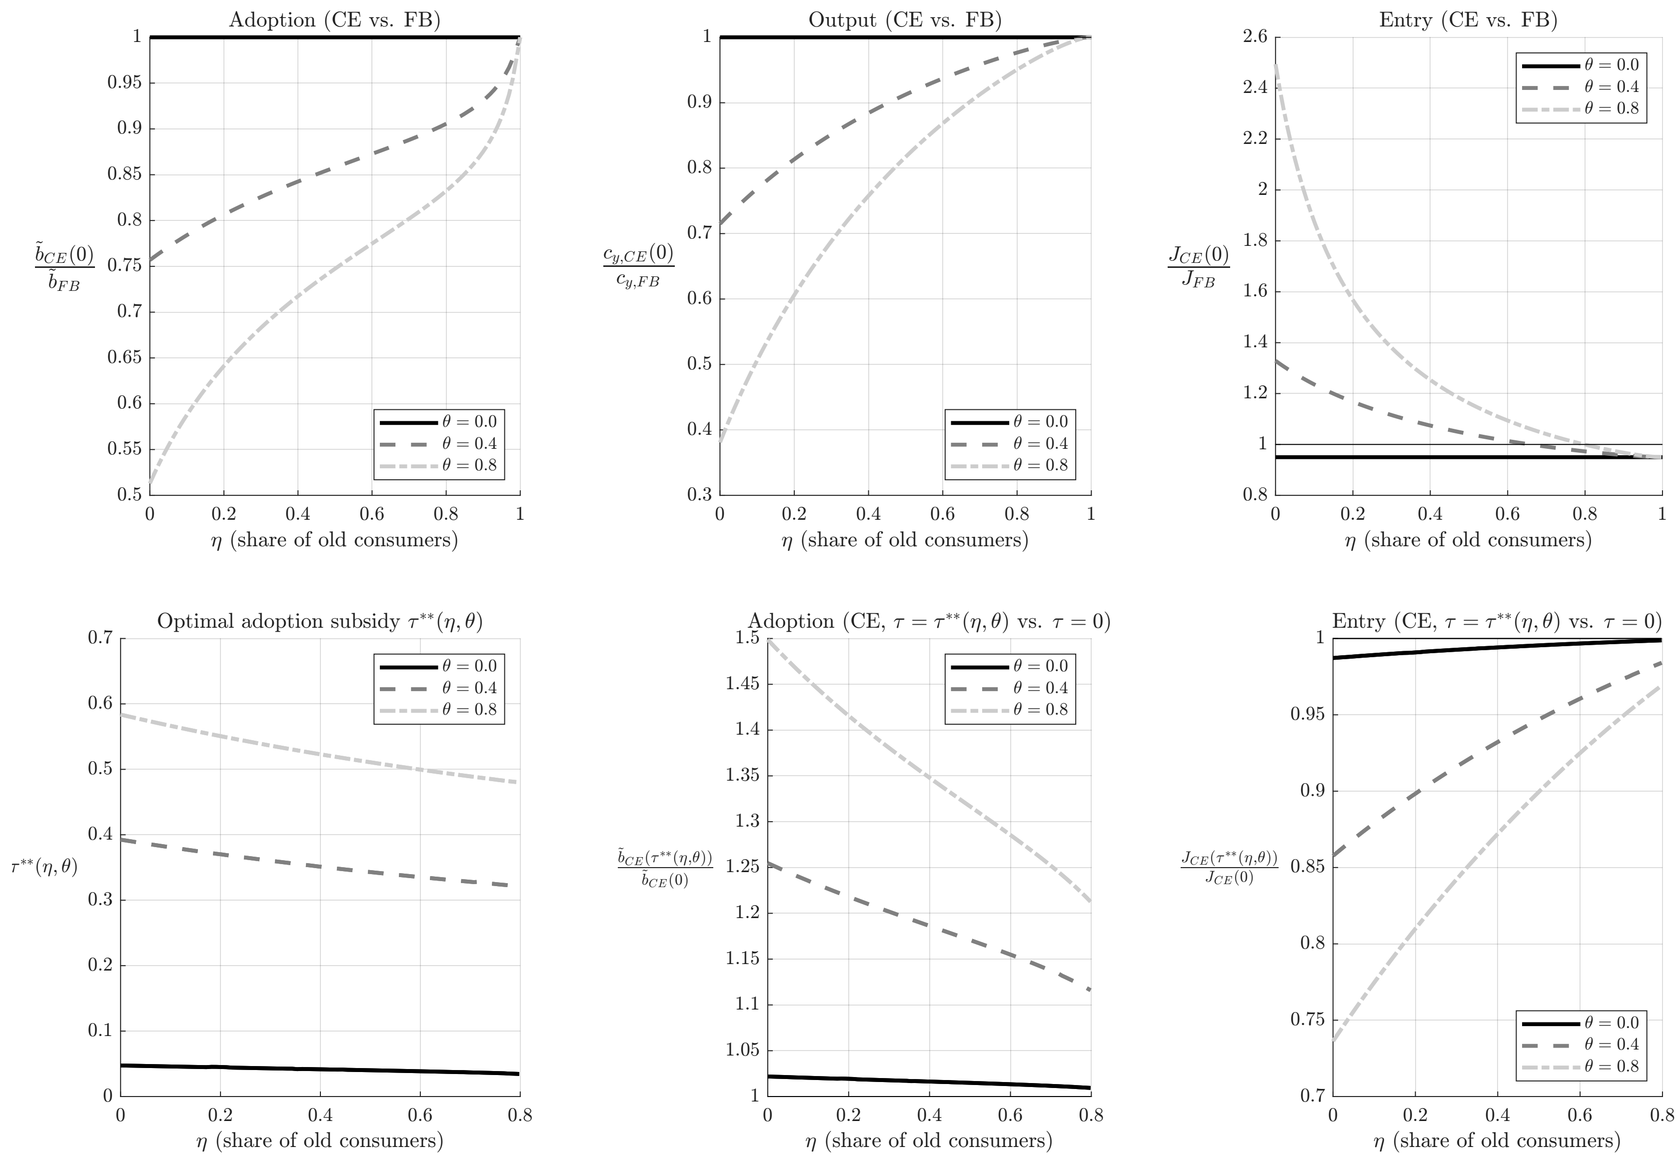

Notes: Competitive equilibrium (CE) and first-best (FB) allocation with endogenous entry. In all figures the horizontal axis is the share η of old consumers. The top row reports comparisons of adoption, production and entry in the CE with no subsidy ($\tau = 0$) and in the FB. The bottom row reports the optimal adoption subsidy (left panel) and a comparison of the CE with optimal subsidy ($\tau = \tau^{**}(\eta, \theta)$) vs. the CE with no subsidy ($\tau = 0$). The calibration used is $E = 1$, $\alpha = 0.9$, $\rho = 0.5$, $\nu = 0.1$, $\xi = 0.1$, and $\gamma(b) = \frac{1}{2}\omega(b-1)^2$ with $\omega = 0.025$. Results are reported for different degrees of network externalities.

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 (Huettner 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)
$\text{AgeStructure}_d \times \text{Post}_t$	0.005***	0.008***	0.083***	0.162***
	[0.002]	[0.002]	[0.028]	[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 16. 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	R^2	Coefficient	R^2
	(1)	(2)	(3)	(4)
Population (IHS)	-0.192***	0.042		
	(0.045)			
Share of agricultural workers	-0.023***	0.048		
	(0.005)			
Number of firms (IHS)	-0.360***	0.116		
	(0.047)			
Literacy Rate	-0.049***	0.222		
	(0.004)			
Share of working population	-0.015***	0.051		
	(0.003)			
Night lights (IHS)	-0.069	0.003		
	(0.063)			
Total stores (IHS)	-0.610***	0.077	-0.056	0.612
	(0.080)		(0.082)	
Total new stores (IHS)	-0.512***	0.065	-0.118	0.562
	(0.070)		(0.072)	
Total transaction volumne (IHS)	-0.784***	0.081	-0.055	0.548
	(0.105)		(0.121)	
Total transaction amount (IHS)	-0.859***	0.055	-0.018	0.417
	(0.144)		(0.189)	
Total rural population (IHS)	0.003	0.000	0.046	0.270
	(0.105)		(0.119)	
Number of schools (IHS)	-0.120***	0.023	0.039	0.781
	(0.044)		(0.035)	
Population density	-0.127***	0.017	-0.075	0.488
	(0.048)		(0.066)	
Bank Branch Density	-0.113*	0.005	-0.010	0.044
	(0.059)		(0.064)	
Share of manufacturing workers	-0.005	0.002	-0.003	0.221
	(0.005)		(0.008)	
Share of small firms	-0.000	0.000	-0.000	0.257
	(0.000)		(0.001)	
Share of primary education	-0.049***	0.112	-0.005	0.808
	(0.007)		(0.004)	
% of urban HH with mobile phones	-0.005	0.001	-0.013	0.053
	(0.006)		(0.009)	
% of urban HH with computers	0.002	0.001	0.006	0.356
	(0.004)		(0.004)	

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	
	$\text{AgeStructure}_d \times \text{Post}_t$	Adoption rate	# Adoptions (IHS)
	(1)	(2)	(3)
$(\text{Sex Ratio})_{d,1991} \times \text{Post}_t$	61.04***		
	(11.71)		
$(\text{Sex Ratio})^2{}_{d,1991} \times \text{Post}_t$	-25.70***		
	(5.332)		
$\text{AgeStructure}_d \times \text{Post}_t$		0.020***	0.256***
		(0.0046)	(0.085)
Observations	7,722	7,722	7,722
SW F-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 16, 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)
$\mathbf{1}(\text{has university})_p \times \text{Post}_t$	0.232***	0.128***	3.987***	2.832***
	[0.051]	[0.030]	[1.112]	[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 18. 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 adoopted 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 Data Appendix

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

A.1.1 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 Febraury 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.[70] 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).

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.[71] 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-

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

[71] See https://microdata.gov.in/nada43/index.php/catalog/156/overview

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 deposits. 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.1.2 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 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.1.3 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.[72] 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.[73] 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 website,[74] 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.[75] Lastly, the list of Institutions of National Importance was found on the Government website updated at least until 2022.[76] 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

[72] https://web.archive.org/web/20230404082827/https://www.ugc.gov.in/oldpdf/Consolidated_CENTRAL_UNIVERSITIES_List.pdf

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

[74] https://deemed.ugc.ac.in/Home/ListOfDeemedToBeUniversity

[75] https://www.bu.edu/globalprograms/files/2019/02/Private-University-Consolidated-List-Private-Universities.pdf

[76] https://www.education.gov.in/institutions-national-importance

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.[77] 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.

A.2 Appendix to Section 3

In what follows, anticipating our analysis of optimal adoption subsidies, we assume that businesses face a constant adoption subsidy τ per unit of numéraire spent on technology adoption costs. Our baseline competitive equilibrium is the one corresponding to $\tau = 0$.

A.2.1 Competitive equilibrium with fixed number of businesses

Utility maximization for each type of household is equivalent to the following demand system:

$$c_o(j) = \left(\frac{p(j)}{P_o}\right)^{-\frac{1}{1-\rho}} C_o \tag{A1}$$

$$C_o = \alpha \frac{I}{P_o} \tag{A2}$$

$$O_o = (1-\alpha) I \tag{A3}$$

[77]See 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

whereas for young consumers, the demand system is:

$$
\begin{aligned}
c_y(j) &= b(j)\left(\frac{p(j)}{P_y}\right)^{-\frac{1}{1-\rho}} C_y && \text{(A4)}\\
C_y &= \alpha\frac{I}{P_y} && \text{(A5)}\\
O_y &= (1-\alpha)I && \text{(A6)}
\end{aligned}
$$

Here, I is income, which is identical across households:

$$
I = E + \int_0^J \pi(j)dj - \tau\int_0^J \gamma(\tilde{b}(j))dj, \tag{A7}
$$

where note that households are taxed lump-sum in order to fund the adoption subsidy. The two prices indices P_o and P_y are given by:

$$
P_o = \left(\int_0^J p(j)^{-\frac{\rho}{1-\rho}}dj\right)^{-\frac{1-\rho}{\rho}}, \quad P_y = \left(\int_0^J b(j)p(j)^{-\frac{\rho}{1-\rho}}dj\right)^{-\frac{1-\rho}{\rho}} \tag{A8}
$$

Profits for each business can be expressed as:

$$
\pi(j) = (p(j)-\xi)\left(\eta\overbrace{\left(\frac{p(j)}{P_o}\right)^{-\frac{1}{1-\rho}} C_o}^{=c_o(j)} + (1-\eta)\underbrace{\bar{b}^{\theta}\,\tilde{b}(j)\left(\frac{p(j)}{P_y}\right)^{-\frac{1}{1-\rho}} C_y}_{=c_y(j)}\right) - (1-\tau)\gamma(\tilde{b}(j)) - \nu \tag{A9}
$$

Profit maximization for each business leads to the following first-order conditions:

$$
\begin{aligned}
p(j) &= \frac{\xi}{\rho}, && \text{(A10)}\\
\left(1-\frac{\xi}{p(j)}\right)(1-\eta)\frac{p(j)c_y(j)}{\tilde{b}(j)} &= (1-\tau)\gamma'(\tilde{b}(j)). && \text{(A11)}
\end{aligned}
$$

Equations (A12)-(A13) imply that any competitive equilibrium is symmetric because the markup is the same across businesses. Thus we omit the index j in what follows.

The two first-order conditions can be rewritten as:

$$
\begin{aligned}
p &= \frac{\xi}{\rho}, && \text{(A12)}\\
\left(1-\frac{\xi}{p}\right)(1-\eta)\frac{pc_y}{\tilde{b}} &= (1-\tau)\gamma'(\tilde{b}). && \text{(A13)}
\end{aligned}
$$

Profits per business are:

$$
\pi = (p-\xi)\left(\eta c_y + (1-\eta)c_o\right) - (1-\tau)\gamma(\tilde{b}) - \nu. \tag{A14}
$$

Household demand systems imply:

$$P_o = J^{-\frac{1-\rho}{\rho}} p \tag{A15}$$

$$c_o = \alpha \left(\frac{p}{P_o}\right)^{-\frac{1}{1-\rho}} \frac{I}{P_o} = \frac{\alpha I}{Jp} \tag{A16}$$

$$P_y = \tilde{b}^{-(1+\theta)\frac{1-\rho}{\rho}} J^{-\frac{1-\rho}{\rho}} p \tag{A17}$$

$$c_y = \alpha \tilde{b}^{1+\theta} \left(\frac{p}{P_y}\right)^{-\frac{1}{1-\rho}} \frac{I}{P_y} = \alpha \tilde{b}^{-(1+\theta)\frac{1-\rho}{\rho}} J^{-\frac{1}{\rho}} \tilde{b}^{(1+\theta)\frac{1-\rho}{\rho}} J^{\frac{1-\rho}{\rho}} \frac{I}{p} = \frac{\alpha I}{Jp} \tag{A18}$$

Thus profits per business are:

$$\pi = (p-\xi)\frac{\alpha I}{Jp} - (1-\tau)\gamma(\tilde{b}) - \nu. \tag{A19}$$

Substituting this into the definition of household income we obtain:

$$I = E + (p-\xi)\,\frac{\alpha I}{p} - J\left((1-\tau)\gamma(\tilde{b}) + \nu\right) - \tau J \gamma(\tilde{b}) = E + (p-\xi)\,\frac{\alpha I}{p} - J\left(\gamma(\tilde{b}) + \nu\right)$$

Solving for income,

$$I = \frac{1}{1 - \frac{p-\xi}{p}\alpha}\left(E - J\left(\gamma(\tilde{b}) + \nu\right)\right) = \frac{1}{1-(1-\rho)\alpha}\left(E - J\left(\gamma(\tilde{b}) + \nu\right)\right) \tag{A20}$$

This gives the rest of the allocation, which we given below for completeness.

$$(1-\eta)\frac{1-\rho}{\rho}\frac{c_y}{\tilde{b}} = \frac{(1-\tau)\gamma'(\tilde{b})}{\xi} \tag{A21}$$

$$p = \frac{\xi}{\rho} \tag{A22}$$

$$\pi = \frac{1}{1-(1-\rho)\alpha}\left((1-\rho)\alpha\frac{E}{J} - (1-(1-(1-\rho)\alpha)\tau)\gamma(\tilde{b}) - \nu\right) \tag{A23}$$

$$I = \frac{1}{1-(1-\rho)\alpha}\left(E - J\left(\gamma(\tilde{b}) + \nu\right)\right) \tag{A24}$$

$$c_y = \frac{\alpha I}{Jp} \tag{A25}$$

$$P_y = \tilde{b}^{-(1+\theta)\frac{1-\rho}{\rho}} J^{-\frac{1-\rho}{\rho}} p \tag{A26}$$

$$C_y = \frac{\alpha I}{P_y} \tag{A27}$$

$$O_y = (1-\alpha)I \tag{A28}$$

$$c_o = \frac{\alpha I}{Jp} = c_y \tag{A29}$$

$$P_o = J^{-\frac{1-\rho}{\rho}} \frac{\xi}{\rho} \tag{A30}$$

$$O_o = (1-\alpha)I = O_y \tag{A31}$$

In particular, solving for the allocation requires solving for the unique value of $\tilde{b}$ such that:

$$(1-\tau)\gamma'(\tilde{b})\tilde{b} = (1-\eta)\frac{(1-\rho)\alpha}{1-(1-\rho)\alpha}\left(\frac{E}{J} - (\gamma(\tilde{b}) + \nu)\right). \tag{A32}$$

Finally, welfare in the competitive equilibrium with a fixed number of businesses is given by:

$$\begin{aligned}
W &= \eta \log(O_o^{1-\alpha}C_o^{\alpha}) + (1-\eta)\log(O_y^{1-\alpha}C_y^{\alpha}) && \text{(A33)}\\
&= \log((1-\alpha)^{1-\alpha}\alpha^{\alpha}) + \log(I) - \alpha(\eta\log(P_o) + (1-\eta)\log(P_y)) && \text{(A34)}\\
&= \log((1-\alpha)^{1-\alpha}\alpha^{\alpha}) + \alpha\frac{1-\rho}{\rho}\log(J) && \text{(A35)}\\
&+ \log(I) - \alpha\log(p) + \alpha(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\log(\tilde{b}). && \text{(A36)}
\end{aligned}$$

A.2.2 Competitive equilibrium with free-entry

The value of J such that profits are zero in the competitive equilibrium characterized above is:

$$J = \frac{(1-\rho)\alpha E}{(1-[1-(1-\rho)\alpha]\,\tau)\gamma(\tilde{b}) + \nu}. \tag{A37}$$

Substituting we find:

$$I = \frac{\nu + (1-\tau)\gamma(\tilde{b})}{\nu + (1-[1-(1-\rho)\alpha]\,\tau)\gamma(\tilde{b})}E. \tag{A38}$$

Solving for the allocation requires finding $\tilde{b}$ such that:

$$(1-\tau)\gamma'(\tilde{b})\tilde{b} = (1-\eta)\left(\nu + (1-\tau)\gamma(\tilde{b})\right). \tag{A39}$$

The rest of the allocation is given by:

$$p = \frac{\xi}{\rho} \tag{A40}$$

$$\pi = 0 \tag{A41}$$

$$I = \frac{\nu + (1-\tau)\gamma(\tilde{b})}{\nu + (1-[1-(1-\rho)\alpha]\,\tau)\gamma(\tilde{b})}E \tag{A42}$$

$$J = \frac{(1-\rho)\alpha E}{\nu + (1 - [1-(1-\rho)\alpha]\,\tau)\gamma(\tilde{b})} \tag{A43}$$

$$\frac{I}{J} = \frac{(1-\tau)\gamma(\tilde{b}) + \nu}{(1-\rho)\alpha} \tag{A44}$$

$$c_y = \frac{\alpha I}{pJ} = \frac{\rho}{1-\rho}\frac{(1-\tau)\gamma(\tilde{b}) + \nu}{\xi} \tag{A45}$$

$$P_y = \tilde{b}^{-(1+\theta)\frac{1-\rho}{\rho}} J^{-\frac{1-\rho}{\rho}} p \tag{A46}$$

$$C_y = \frac{\alpha I}{P_y} \tag{A47}$$

$$O_y = (1-\alpha)I \tag{A48}$$

$$c_o = \frac{\alpha I}{Jp} = c_y \tag{A49}$$

$$P_o = J^{-\frac{1-\rho}{\rho}} \frac{\xi}{\rho} \tag{A50}$$

$$O_o = (1-\alpha)I = O_y \tag{A51}$$

In particular, when $\tau = 0$, we have:

$$\gamma'(\tilde{b})\tilde{b} = (1-\eta)\left(\nu + \gamma(\tilde{b})\right) \tag{A52}$$

$$p = \frac{\xi}{\rho} \tag{A53}$$

$$\pi = 0 \tag{A54}$$

$$I = E \tag{A55}$$

$$J = \frac{(1-\rho)\alpha E}{\gamma(\tilde{b}) + \nu} \tag{A56}$$

$$\frac{I}{J} = \frac{\gamma(\tilde{b}) + \nu}{(1-\rho)\alpha} \tag{A57}$$

$$c_y = \frac{\rho}{1-\rho}\frac{\gamma(\tilde{b}) + \nu}{\xi} \tag{A58}$$

Finally, welfare in the competitive equilibrium with free-entry is given by:

$$W = \log((1-\alpha)^{1-\alpha}\alpha^{\alpha}) + \alpha\frac{1-\rho}{\rho}\log(J) \tag{A59}$$

$$+ \quad \log(I) - \alpha \log(p) + \alpha(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\log(\tilde{b}). \tag{A60}$$

A.2.3 Proofs

Proof of result 1 with fixed number of businesses. When $\eta < 1$, the condition characterizing $\tilde{b}$ when the number of businesses is fixed can be written as:

$$g(\tilde{b};\eta) = \frac{E}{J} - \nu \tag{A61}$$

$$g(\tilde{b};\eta) \equiv \frac{1-\tau}{1-\eta}\frac{1-(1-\rho)\alpha}{(1-\rho)\alpha}\gamma'(\tilde{b})\tilde{b} + \gamma(\tilde{b}) \tag{A62}$$

Differentiating with respect to η,

$$\frac{\partial \tilde{b}}{\partial \eta} = -\frac{\partial g/\partial \eta}{\partial g/\partial \tilde{b}} \tag{A63}$$

Because γ is increasing and convex, $\partial g/\partial \tilde{b} > 0$. Moreover,

$$\frac{\partial g}{\partial \eta} = \frac{1-\tau}{(1-\eta)^2}\frac{1-(1-\rho)\alpha}{(1-\rho)\alpha}\gamma'(\tilde{b})\tilde{b} > 0, \tag{A64}$$

establishing the result. ■

Proof of result 1 with free-entry. When $\eta < 1$, the condition characterizing $\tilde{b}$ when there is free-entry can be written as:

$$g(\tilde{b};\eta) = \nu \tag{A65}$$

$$g(\tilde{b};\eta) \equiv \frac{1-\tau}{1-\eta}\gamma'(\tilde{b})\tilde{b} - (1-\tau)\gamma(\tilde{b}) \tag{A66}$$

Differentiating with respect to η,

$$\frac{\partial \tilde{b}}{\partial \eta} = -\frac{\partial g/\partial \eta}{\partial g/\partial \tilde{b}} \tag{A67}$$

We have:

$$\frac{\partial g}{\partial \eta} = \frac{1-\tau}{(1-\eta)^2}\gamma'(\tilde{b})\tilde{b} > 0. \tag{A68}$$

Moreover,

$$\frac{\partial g}{\partial \tilde{b}} = \left(\frac{1-\tau}{1-\eta} - (1-\tau)\right)\gamma'(\tilde{b}) + \frac{1-\tau}{1-\eta}\tilde{b}\gamma''(\tilde{b}). \tag{A69}$$

When $\tau = 0$, this becomes:

$$\frac{\partial g}{\partial \tilde{b}} = \left(\frac{1}{1-\eta} - 1\right)\gamma'(\tilde{b}) + \frac{1}{1-\eta}\tilde{b}\gamma''(\tilde{b}) > 0 \tag{A70}$$

because $\gamma(.)$ is increasing and convex, establishing the result. ■

Proof of result 2 with fixed number of businesses. When $\eta < 1$, the condition characterizing $\tilde{b}$ when the number of businesses is fixed and $\tau = 0$ can be written as:

$$g(\tilde{b}; \eta, \omega) = \frac{E}{J} - \nu \tag{A71}$$

$$g(\tilde{b}; \eta, \omega) \equiv \omega\left(\frac{1}{1-\eta}\underbrace{\frac{1-(1-\rho)\alpha}{(1-\rho)\alpha}}_{\equiv\sigma} h'(\tilde{b})\tilde{b} + h(\tilde{b})\right) \tag{A72}$$

Differentiating with respect to ω,

$$\frac{\partial \tilde{b}}{\partial \omega} = -\frac{\partial g/\partial \omega}{\partial g/\partial \tilde{b}} \tag{A73}$$

Because h is increasing and convex, $\partial g/\partial \tilde{b} > 0$. Moreover,

$$\frac{\partial g}{\partial \omega} = \frac{g}{\omega} > 0, \tag{A74}$$

establishing that $\frac{\partial \tilde{b}}{\partial \omega} < 0$.

We have:

$$\frac{\partial g}{\partial \tilde{b}}\frac{\partial \tilde{b}}{\partial \omega} + \frac{g}{\omega} = 0 \tag{A75}$$

$$\frac{\partial g}{\partial \tilde{b}}\frac{\partial \tilde{b}}{\partial \eta} + \frac{\partial g}{\partial \eta} = 0 \tag{A76}$$

Differentiating Equation (A75) with respect to η and using Equation (A76) we obtain:

$$\frac{\partial^2 \tilde{b}}{\partial \omega \partial \eta} = -\left(\underbrace{\frac{\partial g}{\partial \tilde{b}}}_{>0}\right)^{-1} \underbrace{\frac{\partial \tilde{b}}{\partial \omega}}_{<0}\left(\frac{\partial^2 g}{\partial \tilde{b}\partial \eta} + \frac{\partial^2 g}{\partial \tilde{b}^2}\frac{\partial \tilde{b}}{\partial \eta}\right) \tag{A77}$$

Next, we note the following:

$$\frac{\partial g}{\partial \tilde{b}} = \omega\left(\frac{1}{1-\eta}\sigma(2\tilde{b}-1) + \tilde{b} - 1\right) \tag{A78}$$

$$\frac{\partial^2 g}{\partial \tilde{b}^2} = \omega\left(\frac{1}{1-\eta}2\sigma+1\right) \tag{A79}$$

$$\frac{\partial g}{\partial \eta} = \omega\frac{1}{(1-\eta)^2}\tilde{b}(\tilde{b}-1) \tag{A80}$$

$$\frac{\partial^2 g}{\partial \eta \partial \tilde{b}} = \omega\frac{1}{(1-\eta)^2}(2\tilde{b}-1) \tag{A81}$$

$$\frac{\partial \tilde{b}}{\partial \eta} = -\frac{\tilde{b}(\tilde{b}-1)}{(1-\eta)\sigma(2\tilde{b}-1)+(1-\eta)^2(\tilde{b}-1)} \tag{A82}$$

$$\frac{\partial^2 g}{\partial \tilde{b}\partial \eta}+\frac{\partial^2 g}{\partial \tilde{b}^2}\frac{\partial \tilde{b}}{\partial \eta} = \omega\frac{1}{(1-\eta)^2}(2\tilde{b}-1)-\omega\left(\frac{1}{1-\eta}2\sigma+1\right)\frac{\tilde{b}(\tilde{b}-1)}{(1-\eta)\sigma(2\tilde{b}-1)+(1-\eta)^2(\tilde{b}-1)}$$

This last term is positive if and only if:

$$\left(2\tilde{b}-1\right)\left((1-\eta)\sigma(2\tilde{b}-1)+(1-\eta)^2(\tilde{b}-1)\right) > \left((1-\eta)2\sigma+(1-\eta)^2\right)\tilde{b}(\tilde{b}-1), \tag{A83}$$

which is equivalent to:

$$(1-\eta)\sigma(2\tilde{b}^2+1)+(1-\eta)^2(\tilde{b}-1)^2 > 0, \tag{A84}$$

which always holds, establishing the result. ■

Proof of result 2 with free-entry. When $\eta < 1$, the condition characterizing $\tilde{b}$ with free-entry and $\tau = 0$ can be written as:

$$g(\tilde{b};\eta,\omega) = \nu \tag{A85}$$

$$g(\tilde{b};\eta,\omega) \equiv \omega\left(\frac{1}{1-\eta}h'(\tilde{b})\tilde{b}+h(\tilde{b})\right) \tag{A86}$$

The derivations are then identical to those above, setting $\sigma = 1$. Since none of the results depend on the value of σ (other than it being positive), this establishes the result. ■

A.3 Appendix to Section 5

A.3.1 The short run: fixed number of businesses

Proof of result 1. The first-best allocation is the solution to:

$$W = \max_{O_o,O_y,c_o,c_y,\tilde{b}} \eta\log(O_o^{1-\alpha}C_o^{\alpha})+(1-\eta)\log(O_y^{1-\alpha}C_y^{\alpha}) \tag{A87}$$

$$\text{s.t.} \quad \eta O_o+(1-\eta)O_y+J\xi\left(\eta c_o+(1-\eta)c_y\right)+J(\gamma(\tilde{b})+\nu) \leq E \qquad [\lambda]$$

$$C_O = J^{\frac{1}{\rho}} c_o \tag{A88}$$

$$C_y = J^{\frac{1}{\rho}} \tilde{b}^{(1+\theta)\frac{1-\rho}{\rho}} c_y \tag{A89}$$

The solution is:

$$\begin{aligned}
\tilde{b}\gamma'(\tilde{b}) &= \alpha(1-\eta)(1+\theta)\frac{(1-\rho)}{\rho}\left(\frac{E}{J} - (\gamma(\tilde{b}) + \nu)\right) && \text{(A90)} \\
\lambda^{-1} &= E - J(\gamma(\tilde{b}) + \nu) \\
O_y &= (1-\alpha)\lambda^{-1} \\
O_o &= O_y \\
c_y &= \alpha\frac{\lambda^{-1}}{J\xi} \\
c_o &= c_y && \text{(A91)}
\end{aligned}$$

Note that in particular, $(c_y, \tilde{b})$ satisfy:

$$(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\frac{c_y}{\tilde{b}} = \frac{\gamma'(\tilde{b})}{\xi}. \tag{A92}$$

Any price system consistent with utility maximization of either young or old household must yield an equilibrium level of income I that satisfies $O_o = O_y = (1-\alpha)I$. Thus:

$$I = \lambda^{-1} = E - J(\gamma(\tilde{b}) + \nu). \tag{A93}$$

Moreover, the budget constraint of young (and old) households must hold with equality. Letting p be the price of each variety we must have:

$$I = \lambda^{-1} = O_y + Jpc_y = \lambda^{-1}\left(1 - \alpha + \alpha\frac{p}{\xi}\right), \tag{A94}$$

so that:

$$p = \xi. \tag{A95}$$

Thus there is a unique price such that the allocation above is consistent with household utility maximization. Either Equation (A93) or (A95) imply that under this price, business profits are:

$$\pi = -(\gamma(\tilde{b}) + \nu). \tag{A96}$$

Finally, like in the competitive equilibrium, welfare can be written as:

$$W = \log((1-\alpha)^{1-\alpha}\alpha^{\alpha}) + \alpha\frac{1-\rho}{\rho}\log(J) \tag{A97}$$

$$+ \log(I) - \alpha\log(p) + \alpha(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\log(\tilde{b}) \tag{A98}$$

for the income I, price p and technology adoption $\tilde{b}$ derived above.

The two conditions determining the optimal level of technology adoption in the competitive equilibrium (when $\tau = 0$) and the first-best can be written as:

$$\tilde{b}\gamma'(\tilde{b}) = Z_{FB}\left(\frac{E}{J} - (\gamma(\tilde{b}) + \nu)\right), \qquad Z_{FB} = \alpha(1-\eta)(1+\theta)\frac{(1-\rho)}{\rho} \tag{A99}$$

$$\tilde{b}\gamma'(\tilde{b}) = Z_{CE}\left(\frac{E}{J} - (\gamma(\tilde{b}) + \nu)\right), \qquad Z_{CE} = (1-\eta)\frac{(1-\rho)\alpha}{1-(1-\rho)\alpha} \tag{A100}$$

Thus a sufficient condition for $\tilde{b}_{CE} < \tilde{b}_{FB}$ is that:

$$Z_{CE} < Z_{FB}, \tag{A101}$$

which is equivalent to:

$$1 + \theta(1-(1-\rho)\alpha) < (1-\rho)\alpha + \rho. \tag{A102}$$

The latter condition is true for any $\theta \geq 0$. ■

Proof of result 2. Recall that welfare in the competitive equilibrium can be written as:

$$W = \log((1-\alpha)^{1-\alpha}\alpha^{\alpha}) + \alpha\frac{1-\rho}{\rho}\log(J) \tag{A103}$$

$$+ \log(I) - \alpha\log(p) + \alpha(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\log(\tilde{b}). \tag{A104}$$

where:

$$(1-\tau)\gamma'(\tilde{b})\tilde{b} = (1-\eta)\frac{(1-\rho)\alpha}{1-(1-\rho)\alpha}\left(\frac{E}{J} - (\gamma(\tilde{b}) + \nu)\right) \tag{A105}$$

$$p = \frac{\xi}{\rho} \tag{A106}$$

$$I = \frac{1}{1-(1-\rho)\alpha}\left(E - J\left(\gamma(\tilde{b}) + \nu\right)\right) \tag{A107}$$

The optimal subsidy is the solution to:

$$\tau^* = \arg\max_{\tau} \quad \log(I) - \alpha\log(p) + \alpha(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\log(\tilde{b}) \tag{A108}$$

subject to Equations (A105)-(A107). The first-order condition is:

$$\frac{\partial\tilde{b}}{\partial\tau}\left(\frac{1}{I}\frac{\partial I}{\partial\tilde{b}} + \alpha(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\frac{1}{\tilde{b}}\right) = 0. \tag{A109}$$

From equation (A105) we see that $\frac{\partial\tilde{b}}{\partial\tau} > 0$. Then, from Equation (A107) we have:

$$\frac{1}{I}\frac{\partial I}{\partial\tilde{b}} = \frac{-J\gamma'(\tilde{b})}{E - J(\gamma(\tilde{b})+\nu)}. \tag{A110}$$

Thus the optimal tax rate must be such that $\tilde{b}$ satisfies:

$$\tilde{b}\gamma'(\tilde{b}) = \alpha(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\left(\frac{E}{J} - (\gamma(\tilde{b})+\nu)\right), \tag{A111}$$

which implies the same level of adoption as in the first-best (Equation A90). Taking ratios of (A111) and (A106), we see that τ must satisfy:

$$1-\tau^* = \frac{1}{\alpha(1+\theta)}\frac{(1-\rho)\alpha}{1-(1-\rho)\alpha}\frac{\rho}{1-\rho} = \frac{1}{1+\theta}\frac{\rho}{1-(1-\rho)\alpha} \in (0,1)\,. \tag{A112}$$

The corresponding level of consumption (under the optimal tax rate) satisfies:

$$(1-\eta)\frac{1-\rho}{\rho}\frac{c_{y,CE}}{\tilde{b}} = \frac{(1-\tau^*)\gamma'(\tilde{b})}{\xi}. \tag{A113}$$

Substituting the expression for τ^* above,

$$(1+\theta)(1-\eta)\frac{1-\rho}{\rho}\frac{c_{y,CE}}{\tilde{b}} = \frac{\rho}{1-(1-\rho)\alpha}\frac{\gamma'(\tilde{b})}{\xi}. \tag{A114}$$

Comparing this with the expression for the first-best level of consumption:

$$(1+\theta)(1-\eta)\frac{1-\rho}{\rho}\frac{c_{y,FB}}{\tilde{b}} = \frac{\gamma'(\tilde{b})}{\xi}, \tag{A115}$$

we see that:

$$\frac{c_{y,CE}}{c_{y,FB}} = \frac{\rho}{1-(1-\rho)\alpha} < 1. \tag{A116}$$

Thus the optimal subsidy does not achieve the first-best level of consumption or output. ■

A.3.2 The medium run: free entry

Proof of result 3. Using Result 1, the first-best is the solution to:

$$W = \max_J \quad \log((1-\alpha)^{1-\alpha}\alpha^\alpha) + \alpha\frac{1-\rho}{\rho}\log(J) \tag{A117}$$

$$+ \quad \log(I) - \alpha\log(p) + \alpha(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\log(\tilde{b}) \tag{A118}$$

where:

$$\begin{aligned}
\tilde{b}\gamma'(\tilde{b}) &= \alpha(1-\eta)(1+\theta)\frac{(1-\rho)}{\rho}\left(\frac{E}{J} - (\gamma(\tilde{b})+\nu)\right) \\
I &= E - J(\gamma(\tilde{b})+\nu) \\
p &= \xi
\end{aligned}$$

The first-order condition is:

$$\alpha\frac{1-\rho}{\rho}\frac{1}{J} + \frac{1}{I}\frac{\partial I}{\partial J} + \frac{\partial\tilde{b}}{\partial J}\left(\frac{1}{I}\frac{\partial I}{\partial\tilde{b}} + \alpha(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\frac{1}{\tilde{b}}\right) = 0. \tag{A119}$$

But note that:

$$\frac{1}{I}\frac{\partial I}{\partial\tilde{b}} = -\alpha\frac{\gamma'(\tilde{b})}{\xi}\frac{1}{c_y}. \tag{A120}$$

Moreover, from result 1, we have that:

$$(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\frac{1}{\tilde{b}} = \frac{\gamma'(\tilde{b})}{\xi}\frac{1}{c_y}. \tag{A121}$$

Thus in Equation (A119), the term in parentheses is equal to zero at any optimal allocation (for fixed J). Thus the first-order condition simplifies to:

$$\alpha\frac{1-\rho}{\rho}\frac{1}{J} + \frac{1}{I}\frac{\partial I}{\partial J} = \alpha\frac{1-\rho}{\rho}\frac{1}{J} - \frac{\gamma(\tilde{b})+\nu}{E - J(\gamma(\tilde{b})+\nu)} = 0. \tag{A122}$$

Thus the optimal $(J, \tilde{b})$ are the unique solution to the system:

$$\alpha\frac{1-\rho}{\rho}\left(\frac{E}{J} - (\gamma(\tilde{b})+\nu)\right) = \gamma(\tilde{b})+\nu \tag{A123}$$

$$\alpha\frac{1-\rho}{\rho}\left(\frac{E}{J} - (\gamma(\tilde{b})+\nu)\right) = \frac{1}{(1-\eta)(1+\theta)}\tilde{b}\gamma'(\tilde{b}) \tag{A124}$$

implying:

$$
\begin{aligned}
(1-\eta)(1+\theta)(\gamma(\tilde{b})+\nu) &= \tilde{b}\gamma'(\tilde{b}) && \text{(A125)}\\
J &= \frac{\alpha(1-\rho)}{\rho+\alpha(1-\rho)}\frac{E}{\gamma(\tilde{b})+\nu} && \text{(A126)}\\
I &= \frac{\rho}{\rho+\alpha(1-\rho)}E && \text{(A127)}\\
p &= \xi && \text{(A128)}\\
\pi &= -(\gamma(\tilde{b})+\nu) && \text{(A129)}\\
O_y &= (1-\alpha)\frac{\alpha(1-\rho)}{\rho+\alpha(1-\rho)}E && \text{(A130)}\\
O_o &= O_y && \text{(A131)}\\
c_y &= \frac{1}{\xi}\frac{\rho}{1-\rho}(\gamma(\tilde{b})+\nu) && \text{(A132)}
\end{aligned}
$$

Recall that the competitive equilibrium with free-entry is described by:

$$
\begin{aligned}
(1-\eta)(\nu+\gamma(\tilde{b}_{CE})) &= \tilde{b}_{CE}\gamma'(\tilde{b}_{CE}) && \text{(A133)}\\
J_{CE} &= \frac{\alpha(1-\rho)}{\gamma(\tilde{b}_{CE})+\nu}E && \text{(A134)}\\
c_{y,CE} &= \frac{1}{\xi}\frac{\rho}{1-\rho}\left(\gamma(\tilde{b}_{CE})+\nu\right) && \text{(A135)}
\end{aligned}
$$

First consider the case with no externalities: $\theta = 0$. Then we see that $\tilde{b}_{CE} = \tilde{b}_{FB}$, and so $c_{y,CE} = c_{y,FB}$. However, there is still too little entry relative to first-best: $J_{CE} < J_{FB}$.

Now consider the case with externalities: $\theta > 0$. Because $(1-\eta)(1+\theta)(\gamma(\tilde{b})+\nu) > (1-\eta)(\gamma(\tilde{b})+\nu)$, there is too little adoption in the competitive equilibrium:

$$\tilde{b}_{CE} < \tilde{b}_{FB}. \tag{A136}$$

This also implies that output is too low in the equilibrium with free-entry, $c_{y,CE} < c_{y,FB}$. Finally, whether entry is higher in CE or FB is ambiguous, and depends on whether:

$$(\alpha(1-\rho)+\rho)(\gamma(\tilde{b}_{FB})+\nu) > \gamma(\tilde{b}_{CE})+\nu. \tag{A137}$$

From Equation (A125), one sees that $\lim_{\theta\to+\infty} b_{FB}(\theta) = +\infty$, while $\tilde{b}_{CE}$ is independent of θ; thus for sufficiently large θ the condition above will hold and entry in the competitive equilibrium will be too high relative to that required by the planner. ■

proof of result 4. Recall that welfare in the competitive equilibrium can be written as:

$$W = \log((1-\alpha)^{1-\alpha}\alpha^{\alpha}) + \alpha\frac{1-\rho}{\rho}\log(J) \tag{A138}$$

$$+ \log(I) - \alpha\log(p) + \alpha(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\log(\tilde{b}). \tag{A139}$$

where:

$$(1-\tau)\gamma'(\tilde{b})\tilde{b} = (1-\eta)\left(\nu + (1-\tau)\gamma(\tilde{b})\right) \tag{A140}$$

$$p = \frac{\xi}{\rho} \tag{A141}$$

$$I = \frac{\nu + (1-\tau)\gamma(\tilde{b})}{\nu + (1-[1-(1-\rho)\alpha]\,\tau)\gamma(\tilde{b})}E \tag{A142}$$

$$J = \frac{(1-\rho)\alpha E}{\nu + (1-[1-(1-\rho)\alpha]\,\tau)\gamma(\tilde{b})} \tag{A143}$$

The optimal subsidy is the solution to:

$$\tau^{**} = \arg\max_{\tau} \quad \alpha\frac{1-\rho}{\rho}\log(J) + \log(I) + \alpha(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\log(\tilde{b}) \tag{A144}$$

subject to Equations (A140), (A142) and (A143).

We first rewrite the objective as:

$$W(\tau;\eta,\theta) = \left(1 + \alpha\frac{1-\rho}{\rho}\right)\log(J) + \log(\Omega) + \alpha(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\log(\tilde{b}) \tag{A145}$$

where:

$$(1-\tau)\gamma'(\tilde{b})\tilde{b} = (1-\eta)\left(\nu + (1-\tau)\gamma(\tilde{b})\right) \tag{A146}$$

$$\frac{I}{J} \equiv \Omega = \frac{\nu + (1-\tau)\gamma(\tilde{b})}{(1-\rho)\alpha} \tag{A147}$$

$$J = \frac{(1-\rho)\alpha E}{\nu + (1-[1-(1-\rho)\alpha]\,\tau)\gamma(\tilde{b})} \tag{A148}$$

Next, we have the following relationships:

$$\frac{\partial J}{\partial \tau} = (1-(1-\rho)\alpha)\gamma(\tilde{b})\frac{J^2}{(1-\rho)\alpha E} \tag{A149}$$

$$\frac{\partial J}{\partial \tilde{b}} = -(1-[1-(1-\rho)\alpha]\,\tau)\gamma'(\tilde{b})\frac{J^2}{(1-\rho)\alpha E} \tag{A150}$$

$$\frac{\partial \Omega}{\partial \tau} = \frac{-\gamma(\tilde{b})}{\nu+(1-\tau)\gamma(\tilde{b})}\Omega \tag{A151}$$

$$\frac{\partial \Omega}{\partial \tilde{b}} = \frac{-\gamma(\tilde{b})}{\nu+(1-\tau)\gamma'(\tilde{b})}\Omega \tag{A152}$$

The necessary first-order condition characterizing τ^{**} can be written as:

$$\begin{aligned} 0=\frac{\partial W}{\partial \tau}(\tau;\eta,\theta) &= \left(1+\alpha\frac{1-\rho}{\rho}\right)\frac{1}{J}\left(\frac{\partial J}{\partial \tau}+\frac{\partial J}{\partial \tilde{b}}\frac{\partial \tilde{b}}{\partial \tau}\right) \\ &+ \frac{1}{\Omega}\left(\frac{\partial \Omega}{\partial \tau}+\frac{\partial \Omega}{\partial \tilde{b}}\frac{\partial \tilde{b}}{\partial \tau}\right) \\ &+ \alpha(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\frac{1}{\tilde{b}}\frac{\partial \tilde{b}}{\partial \tau} \end{aligned}$$

Using the equations above, we can reorganize the expression for the derivative of the objective function as:

$$\begin{aligned} \frac{\partial W}{\partial \tau}(\tau;\eta,\theta) &= \frac{\gamma(\tilde{b})}{\nu+(1-\tau)\gamma(\tilde{b})}\left[\left(1+\alpha\frac{1-\rho}{\rho}\right)(1-(1-\rho)\alpha)\frac{\nu+(1-\tau)\gamma(\tilde{b})}{\nu+(1-(1-(1-\rho)\alpha)\,\tau)\gamma(\tilde{b})}-1\right] \\ &+ (1-\eta)\frac{1}{\tilde{b}}\frac{\partial \tilde{b}}{\partial \tau}\left[1+(1+\theta)\alpha\frac{1-\rho}{\rho}-\left(1+\alpha\frac{1-\rho}{\rho}\right)\frac{1-(1-(1-\rho)\alpha)\tau}{1-\tau}\frac{\nu+(1-\tau)\gamma(\tilde{b})}{\nu+(1-(1-(1-\rho)\alpha)\,\tau)\gamma(\tilde{b})}\right] \end{aligned}$$

Finally, note that:

$$\frac{\nu+(1-\tau)\gamma(\tilde{b})}{\nu+(1-(1-(1-\rho)\alpha)\,\tau)\gamma(\tilde{b})}=\frac{I}{E}, \quad \geq 1, \quad >1 \quad \text{iff } \tau>0.$$

So we write the derivative of the objective function in more condensed form as:

$$\begin{aligned} \frac{\partial W}{\partial \tau}(\tau;\eta,\theta) &= \frac{\gamma(\tilde{b})}{\nu+(1-\tau)\gamma(\tilde{b})}\left[\left(1+\alpha\frac{1-\rho}{\rho}\right)(1-(1-\rho)\alpha)\frac{I}{E}-1\right] \\ &+ (1-\eta)\frac{1}{\tilde{b}}\frac{\partial \tilde{b}}{\partial \tau}\left[1+(1+\theta)\alpha\frac{1-\rho}{\rho}-\left(1+\alpha\frac{1-\rho}{\rho}\right)\frac{1-(1-(1-\rho)\alpha)\tau}{1-\tau}\frac{I}{E}\right] \end{aligned}$$

Next, note that:

$$\frac{\partial \tilde{b}}{\partial \tau} = \frac{1}{(1-\tau)^2}\frac{(1-\eta)\nu}{\gamma''(\tilde{b})\tilde{b}+\eta\gamma'(\tilde{b})}>0$$

$$\frac{\partial\tilde{b}}{\partial\eta} = \frac{1}{1-\tau}\frac{-(\nu+(1-\tau)\gamma(\tilde{b}))}{\gamma''(\tilde{b})\tilde{b}+\eta\gamma'(\tilde{b})} < 0$$

Now we prove that $\tau^{**} > 0$, even when $\theta = 0$. We have:

$$\begin{aligned}\frac{\partial W}{\partial\tau}(0;\eta,\theta) &= \frac{\gamma(\tilde{b})}{\nu+\gamma(\tilde{b})}\left[\left(1+\alpha\frac{1-\rho}{\rho}\right)(1-(1-\rho)\alpha)-1\right] \\ &+ (1-\eta)\frac{1}{\tilde{b}}\frac{\partial\tilde{b}}{\partial\tau}\theta\alpha\frac{1-\rho}{\rho}\end{aligned}$$

Note that:

$$\left(1+\alpha\frac{1-\rho}{\rho}\right)(1-(1-\rho)\alpha) = 1+(1-\alpha)(1-\rho)\alpha\frac{1-\rho}{\rho}$$

So:

$$\begin{aligned}\frac{\partial W}{\partial\tau}(0;\eta,\theta) &= \alpha\frac{1-\rho}{\rho}\left(\frac{\gamma(\tilde{b})}{\nu+\gamma(\tilde{b})}(1-\alpha)(1-\rho)+(1-\eta)\frac{1}{\tilde{b}}\frac{\partial\tilde{b}}{\partial\tau}\theta\right) \\ &> 0,\end{aligned}$$

so that it must be that $\tau^{**} > 0$. This is true even when $\theta = 0$.

Next we prove that when $\theta = 0$, the welfare gain from a small adoption subsidy is lower, the higher the share of old consumers. We have:

$$\begin{aligned}\frac{\partial}{\partial\eta}\frac{\partial W}{\partial\tau}(0;\eta,\theta) &\propto -\frac{\theta}{\tilde{b}}\frac{\partial\tilde{b}}{\partial\tau} \\ &+ (1-\eta)\theta\frac{\partial}{\partial\eta}\left(\frac{1}{\tilde{b}}\frac{\partial\tilde{b}}{\partial\tau}\right) \\ &+ (1-\alpha)(1-\rho)\frac{\partial}{\partial\eta}\left(1-\frac{\nu}{\nu+\gamma(\tilde{b})}\right)\end{aligned}$$

The last term is always negative because $\partial\tilde{b}/\partial\eta < 0$. Additionally, when $\theta = 0$, the first two terms are equal to zero, establishing the result. ■

A.3.3 Comparing the CE and the constrained optimal (CO) allocation

We use "CO" to denote the planner's problem that maximizes welfare *subject to* the same technological and market constraints as the competitive equilibrium, most importantly, positive profits. This benchmark is distinct from the first best (FB), which can set $p = \xi$ and transfer business losses to the household. Under free entry, with $\theta = 0$, CO and CE coincide (the classic constrained-efficiency result), whereas with a fixed number of businesses CO can differ from CE.

1. The short-run: fixed number of businesses

Definition 5. *A constrained optimal (CO) allocation is a price p for intermediate varieties and a set of quantities $\left\{O_o, O_y, c_o, c_y, \tilde{b}\right\}$ that maximizes the welfare criterion (19) subject to the constraints that (a) the quantities $\{O_o, O_y, c_o, c_y\}$ are consistent with utility maximization of young and old consumers given $(p, \tilde{b})$; (b) resulting business profits are weakly positive, $\pi \geq 0$; and (c) the resource constraint (20) holds.*

Result 5 (Constrained-optimal allocation)**.** *In the constrained optimal allocation, the planner also chooses $(c_y, \tilde{b})$ such that (21) holds. Price is equal to average cost and businesses make zero profits. Finally, when profits in the competitive equilibrium are strictly positive, technology adoption and consumption are strictly higher than in the CE, even without externalities ($\theta = 0$):*

$$\tilde{b}_{CO} > \tilde{b}_{CE}, \quad c_{y,CO} > c_{y,CE} \quad \text{if} \quad \pi_{CE} > 0. \tag{A153}$$

Proof of result 5. The constrained optimal allocation is the solution to:

$$W = \max_{O_o, O_y, c_o, c_y, \tilde{b}, p} \eta \log(O_o^{1-\alpha} C_o^{\alpha}) + (1-\eta) \log(O_y^{1-\alpha} C_y^{\alpha}) \tag{A154}$$

$$\text{s.t.} \quad \eta O_o + (1-\eta) O_y + J\xi\left(\eta c_o + (1-\eta) c_y\right) + J(\gamma(\tilde{b}) + \nu) \leq E \qquad [\lambda]$$

$$C_O = J^{\frac{1}{\rho}} c_o \tag{A155}$$

$$C_y = J^{\frac{1}{\rho}} \tilde{b}^{(1+\theta)\frac{1-\rho}{\rho}} c_y \tag{A156}$$

$$\pi = (p - \xi) c_y - (\gamma(\tilde{b}) + \nu) \geq 0 \tag{A157}$$

$$c_y = \frac{\alpha(E + J\pi)}{Jp} \tag{A158}$$

Note here that implementability (that is, maximization of utility by consumers given p) requires $O_o = O_y = (1-\alpha)(E + J\pi)$ and $c_o = c_y$, but we have omitted these conditions for brevity, though they are implicit in the expression of business profits.

Ignoring the constraint (A157), we see that the price p only appears in the implementability condition (A158). So we can solve for the optimal allocation ignoring (A158), and use that condition to back out the resulting equilibrium price. But in this case, the problem is identical to the first-best allocation problem discussed above. And this problem, along with the implementability condition (A158), leads to $\pi = -(\gamma(\tilde{b}) + \nu) < 0$, violating the zero-profit condition (A157). Thus, that condition must bind. As a result, at the constrained optimal allocation businesses make zero profits and therefore:

$$I = E, \quad c_o = c_y = \alpha \frac{E}{Jp}, \quad O_o = O_y = (1-\alpha)E. \tag{A159}$$

Substituting these expressions in the definition of profits implies that, for any level of $\tilde{b}$ we must have:

$$\begin{aligned} p &= \xi\left(1+\frac{\gamma(\tilde{b})+\nu}{\alpha\frac{E}{J}-(\gamma(\tilde{b})+\nu)}\right) \\ c_y &= \frac{1}{\xi}\left(\alpha\frac{E}{J}-(\gamma(\tilde{b})+\nu)\right) \end{aligned}$$

Finally, substituting the expression for c_y in the objective function and optimizing with respect to $\tilde{b}$ gives the following solution:

$$\begin{aligned} \tilde{b}\gamma'(\tilde{b}) &= (1-\eta)(1+\theta)\frac{(1-\rho)}{\rho}\left(\alpha\frac{E}{J}-(\gamma(\tilde{b})+\nu)\right) \\ \lambda^{-1} &= E \\ O_y &= (1-\alpha)E \\ O_o &= O_y \\ c_y &= \frac{1}{\xi}\left(\alpha\frac{E}{J}-(\gamma(\tilde{b})+\nu)\right) \\ c_o &= c_y \\ I &= E \\ \pi &= 0 \\ p &= \xi\left(1+\frac{\gamma(\tilde{b})+\nu}{\alpha\frac{E}{J}-(\gamma(\tilde{b})+\nu)}\right) \end{aligned}$$

Note that in particular, $(c_y,\tilde{b})$ satisfy:

$$(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\frac{c_y}{\tilde{b}}=\frac{\gamma'(\tilde{b})}{\xi}. \tag{A160}$$

Finally, like in the competitive equilibrium, welfare can be written as:

$$W = \log((1-\alpha)^{1-\alpha}\alpha^{\alpha})+\alpha\frac{1-\rho}{\rho}\log(J) \tag{A161}$$

$$+ \log(I)-\alpha\log(p)+\alpha(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\log(\tilde{b}) \tag{A162}$$

for the income I, adoption $\tilde{b}$ and price p derived above.

The two conditions determining the optimal level of technology adoption in the constrained optimal allocation and the first-best can be written as:

$$\tilde{b}\gamma'(\tilde{b}) = Z_{FB}\left(\frac{E}{J} - (\gamma(\tilde{b}) + \nu)\right) \qquad (FB) \tag{A163}$$

$$\tilde{b}\gamma'(\tilde{b}) = Z_{FB}\left(\frac{E}{J} - \frac{1}{\alpha}(\gamma(\tilde{b}) + \nu)\right) \qquad (CO) \tag{A164}$$

The right-hand sides of these equations satisfy $Z_{FB}\left(\frac{E}{J} - \frac{1}{\alpha}(\gamma(\tilde{b}) + \nu)\right) \leq Z_{FB}\left(\frac{E}{J} - (\gamma(\tilde{b}) + \nu)\right)$ for all $\tilde{b}$, implying that $\tilde{b}_{FB} > \tilde{b}_{CE}$, and therefore (from Equation (A160), which also holds in the first-best allocation), that $c_{y,FB} > c_{y,CO}$.

Next, we prove that for all $\theta \geq 0$, if profits strictly positive in the CE:

$$(1-\rho)\alpha\frac{E}{J} - (\gamma(\tilde{b}_{CE}) + \nu) > 0, \tag{A165}$$

then $\tilde{b}_{CO} > \tilde{b}_{CE}$. To establish this we proceed in two steps. First we establish it for the case $\theta = 0$. Then we deal with the case $\theta = 0$.

In the case $\theta = 0$, note that both in the CO and CE the following condition links consumption and technology adoption:

$$(1-\eta)\frac{1-\rho}{\rho}\frac{c_y}{\tilde{b}} = \frac{\gamma'(\tilde{b})}{\xi}. \tag{A166}$$

From the solutions derived above, we can write:

$$\begin{aligned} c_{y,CO} - c_{y,CE} &= \frac{1}{\xi}\alpha\left(1 - \frac{\rho}{1-\alpha(1-\rho)}\right)\frac{E}{J} + \frac{1}{\xi}\left(\frac{\alpha\rho}{1-\alpha(1-\rho)}\left(\nu + \gamma(\tilde{b}_{CE})\right) - \left(\nu + \gamma(\tilde{b}_{CO})\right)\right) \\ &= \frac{1}{\xi}\frac{1-\alpha}{1-\alpha(1-\rho)}\left(\alpha(1-\rho)\frac{E}{J} - \left(\nu + \gamma(\tilde{b}_{CE})\right)\right) + \frac{1}{\xi}\left(\gamma(\tilde{b}_{CE}) - \gamma(\tilde{b}_{CO})\right) \end{aligned} \tag{A167}$$

Preparing for a contradiction, assume that:

$$c_{y,CO} \leq c_{y,CE}.$$

Then Equation (A167) implies that we must have $\gamma(\tilde{b}_{CE}) < \gamma(\tilde{b}_{CO})$, since profits are positive in the CE. Therefore, we must also have $\tilde{b}_{CE} < \tilde{b}_{CO}$. But by convexity of γ, Equation (A166 implies that $c_{y,CE} < c_{y,CO}$, a contradiction. Therefore, it must be that:

$$c_{y,CO} > c_{y,CE}.$$

From the first-order conditions with respect to $\tilde{b}$, we then also must have:

$$\tilde{b}_{CO} > \tilde{b}_{CE}.$$

Next, turn to the case $\theta > 0$. Note that $\tilde{b}_{CE}$ is independent of θ. Moreover from the condition determining $\tilde{b}_{CO}$,

$$\tilde{b}\gamma'(\tilde{b}) = (1-\eta)(1+\theta)\frac{(1-\rho)}{\rho}\left(\alpha\frac{E}{J} - (\gamma(\tilde{b})+\nu)\right),$$

it follows that $\frac{\partial \tilde{b}_{CE}}{\partial \theta} > 0$. Thus for any $\theta > 0$,

$$\tilde{b}_{CO}(\theta) > \tilde{b}_{CO}(0) \geq \tilde{b}_{CE}, \tag{A168}$$

establishing the result. ■

2. The long-run: endogenous entry

Definition 6. *A constrained optimal (CO) allocation is a number of businesses J, a price p for intermediate varieties and a set of quantities $\left\{O_o, O_y, c_o, c_y, \tilde{b}\right\}$ that maximizes the welfare criterion (19) subject to the constraints that (a) the quantities $\{O_o, O_y, c_o, c_y\}$ are consistent with utility maximization of young and old consumers given $(p, \tilde{b})$; (b) resulting business profits are weakly positive, $\pi \geq 0$; and (c) the resource constraint (20) holds.*

Result 6 (Constrained optimal allocation)**.** *In the constrained-optimal allocation, the planner chooses a level of adoption that coincides with the first-best and is given by Equation (33). Moreover, the planner chooses prices that are equal to a constant markup over marginal costs, as in the competitive equilibrium. If there are no externalities ($\theta = 0$), technology adoption, consumption and entry coincide with their competitive equilibrium level. With externalities ($\theta > 0$), the constrained-optimal allocation features more technology adoption, more production, and less entry than the competitive equilibrium:*

$$\tilde{b}_{CO} > \tilde{b}_{CE}, \quad c_{y,CO} = c_{o,CO} > c_{y,CE} = c_{o,FB}, \quad J_{CO} < J_{CE}. \tag{A169}$$

Proof of result 6. Using Result 5, the constrained optimum is the solution to:

$$W = \max_J \quad \log((1-\alpha)^{1-\alpha}\alpha^{\alpha}) + \alpha\frac{1-\rho}{\rho}\log(J) \tag{A170}$$

$$+ \quad \log(I) - \alpha\log(p) + \alpha(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\log(\tilde{b}) \tag{A171}$$

where:

$$\begin{aligned}
\tilde{b}\gamma'(\tilde{b}) &= (1-\eta)(1+\theta)\frac{(1-\rho)}{\rho}\left(\alpha\frac{E}{J} - (\gamma(\tilde{b})+\nu)\right) \\
I &= E \\
p &= \xi\left(1 + \frac{\gamma(\tilde{b})+\nu}{\alpha\frac{E}{J} - (\gamma(\tilde{b})+\nu)}\right)
\end{aligned}$$

$$\begin{aligned}
\tilde{b}\gamma'(\tilde{b}) &= (1-\eta)(1+\theta)\frac{(1-\rho)}{\rho}\left(\alpha\frac{E}{J}-(\gamma(\tilde{b})+\nu)\right) \\
I &= E \\
p &= \xi\left(1+\frac{\gamma(\tilde{b})+\nu}{\alpha\frac{E}{J}-(\gamma(\tilde{b})+\nu)}\right)
\end{aligned}$$

The first-order condition is:

$$\alpha\frac{1-\rho}{\rho}\frac{1}{J}-\alpha\frac{1}{p}\frac{\partial p}{\partial J}+\frac{\partial\tilde{b}}{\partial J}\left(-\alpha\frac{1}{p}\frac{\partial p}{\partial\tilde{b}}+\alpha(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\frac{1}{\tilde{b}}\right)=0. \tag{A172}$$

But note that:

$$\frac{1}{p}\frac{\partial p}{\partial\tilde{b}}=\frac{\gamma'(\tilde{b})}{\xi}\frac{1}{c_y}. \tag{A173}$$

Moreover, from result 5, we have that:

$$(1-\eta)(1+\theta)\frac{1-\rho}{\rho}\frac{1}{\tilde{b}}=\frac{\gamma'(\tilde{b})}{\xi}\frac{1}{c_y}. \tag{A174}$$

Thus in Equation (A172), the term in parentheses is equal to zero at any optimal allocation (for fixed J). Thus the first-order condition simplifies to:

$$\frac{1-\rho}{\rho}\frac{1}{J}-\frac{1}{p}\frac{\partial p}{\partial J}=0. \tag{A175}$$

We have:

$$\frac{1}{p}\frac{\partial p}{\partial J}=\frac{1}{\xi}\frac{\gamma(\tilde{b})+\nu}{J}\frac{1}{c_y} \tag{A176}$$

Thus the optimal $(J,\tilde{b})$ are the unique solution to the system:

$$\frac{1-\rho}{\rho}\left(\alpha\frac{E}{J}-(\gamma(\tilde{b})+\nu)\right) = \gamma(\tilde{b})+\nu \tag{A177}$$

$$\frac{1-\rho}{\rho}\left(\alpha\frac{E}{J}-(\gamma(\tilde{b})+\nu)\right) = \frac{1}{(1-\eta)(1+\theta)}\tilde{b}\gamma'(\tilde{b}) \tag{A178}$$

implying:

$$(1-\eta)(1+\theta)(\gamma(\tilde{b})+\nu) = \tilde{b}\gamma'(\tilde{b}) \tag{A179}$$

$$J = \alpha(1-\rho)\frac{E}{\gamma(\tilde{b})+\nu} \tag{A180}$$

$$I = E \tag{A181}$$

$$
\begin{aligned}
p &= \frac{\xi}{\rho} && \text{(A182)} \\
\pi &= 0 && \text{(A183)} \\
O_y &= (1-\alpha)E && \text{(A184)} \\
O_o &= O_y && \text{(A185)} \\
c_y &= \frac{1}{\xi}\frac{\rho}{1-\rho}(\gamma(\tilde{b})+\nu) && \text{(A186)}
\end{aligned}
$$

To compare with the competitive equilibrium with free-entry, note that when $\theta = 0$, the two coincide exactly, so the competitive equilibrium is constrained-efficient. On the other hand, when $\theta > 0$, the constrained optimum features a higher rate of adoption. Thus it also features fewer businesses and a higher rate of output.

Finally, note that when $\theta = 0$, adoption and output are the same in the CO and in the FB, though entry is higher in the FB. When $\theta > 0$, adoption and output are also the same in the CO and in the FB, but entry is strictly lower in the CO than in the FB. ■

Appendix Figures and Tables

Figure A-1: Comparison of distributions across datasets

(a) Distribution across Age

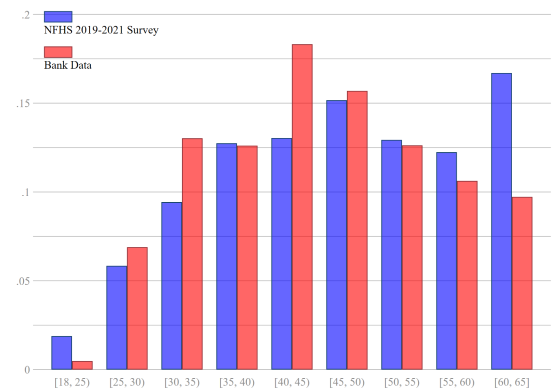

(b) Distribution across income levels

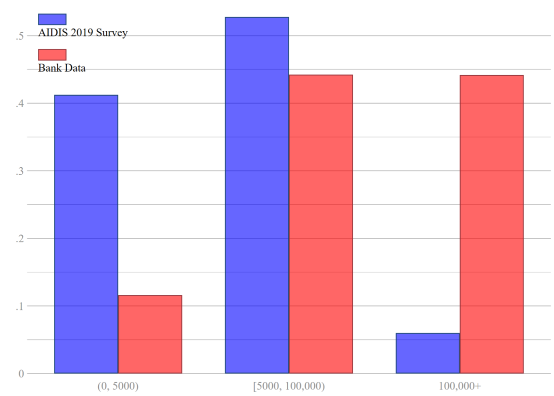

Notes: This figure compares the age and income distribution in our bank level data to the same information provided in nationally representative surveys. Panel (a) reports the age distribution, dividing the sample in 5-year intervals (with the exception of the youngest group that goes from 18 to 25). 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

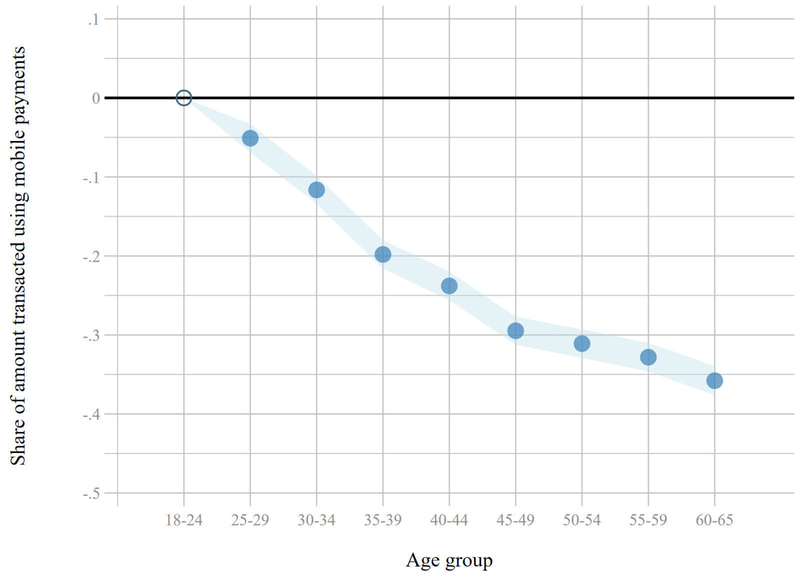

Notes: The figure plots the estimates of the share of the amount transacted using mobile payments by households across different age-groups. We normalize the age groups of 18 to 24 to be zero. 95% confidence intervals are denoted using the blue shared region.

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

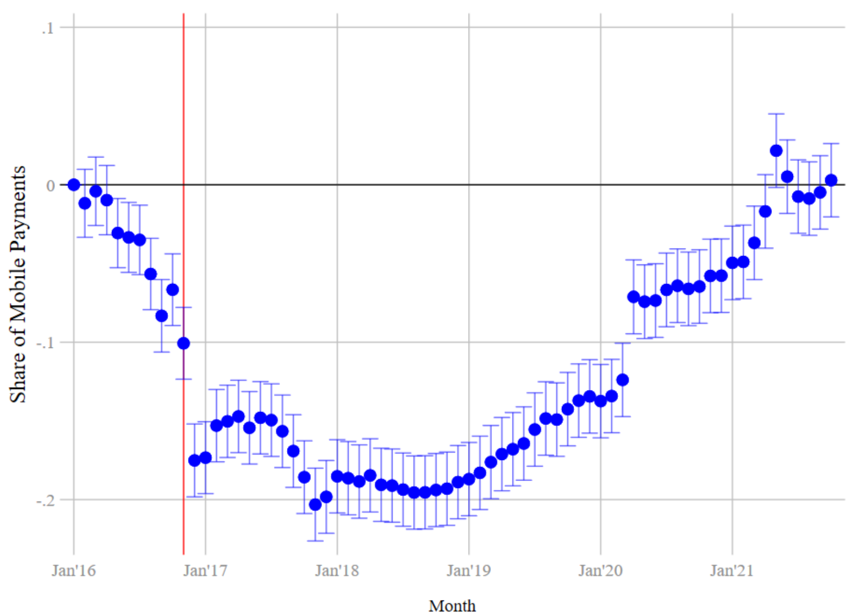

Notes: The figure plots month-by-month the difference in the share of mobile transactions between consumers that were early users of credit cards (i.e., had credit card in 2015) and not early users. The share of mobile is defined as usual. The vertical red-line identifies November 2016 — the month of 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

(a) No controls

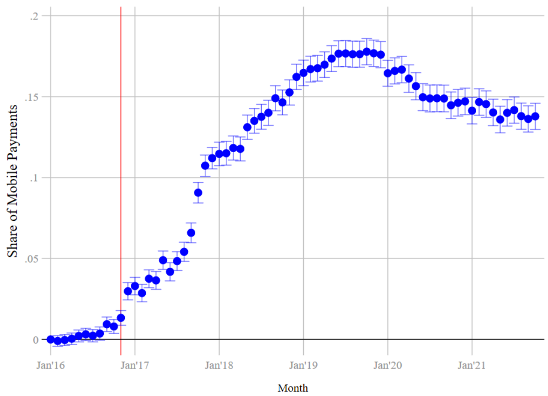

(b) card ownership-month f.e.+ pincode-month f.e.

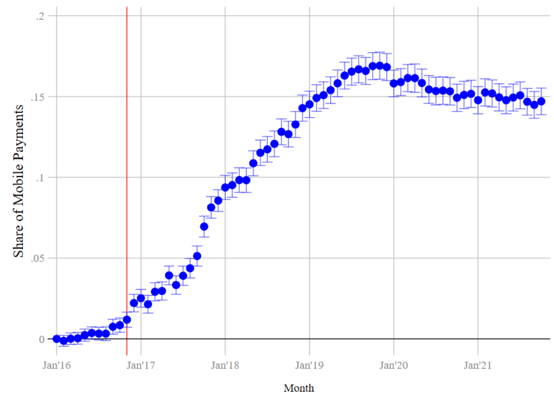

(c) Only card holders

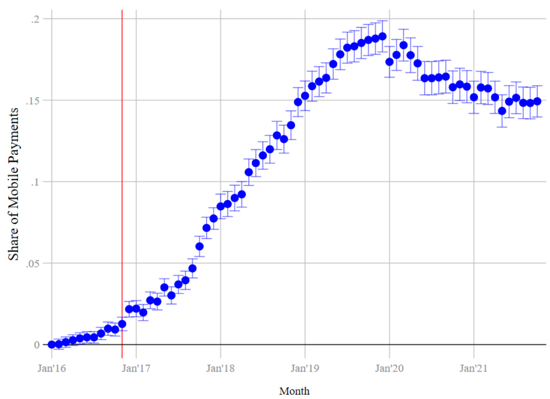

Notes: The figure plots month-by-month the difference in the share of mobile transactions between young and old consumers using the panel from the the bank-level transaction data. We divide the sample across young and old splitting around age 40, defined in 2015. The share of mobile is defined as usual. Panel (a) reports the difference with no controls. Panel (b) reports the difference while also controlling for card ownership-month fixed-effects and pincode-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

(a) baseline specification

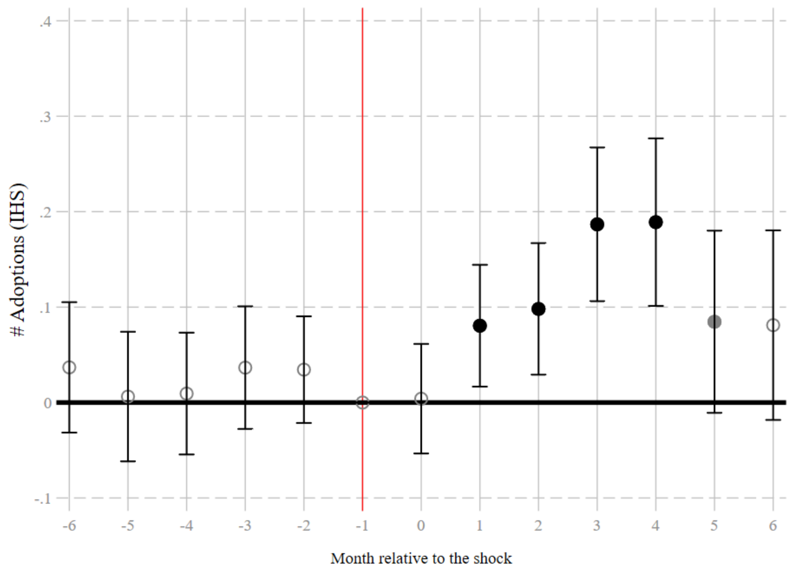

(b) all controls

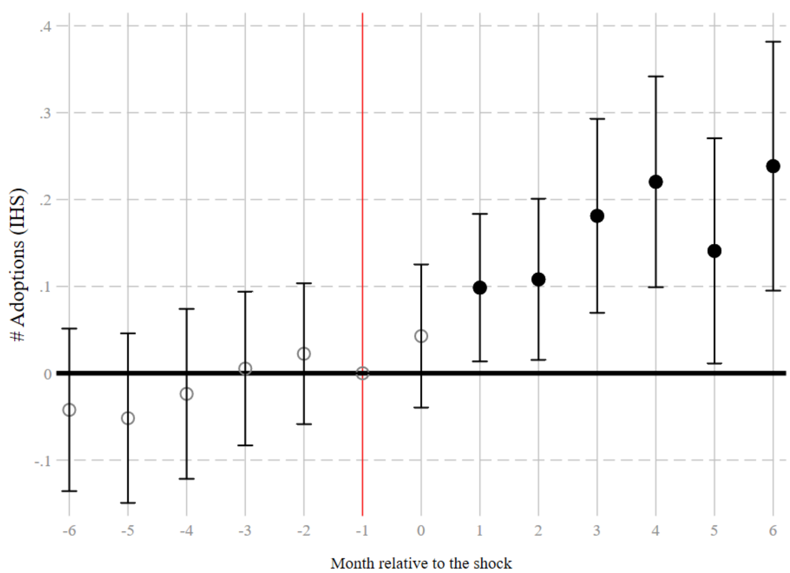

Notes: The figure plots the dynamic treatment effects of age structure on adoption. The dependent variable is the inverse hyperbolic sine (IHS) transformation of number of stores that adopted our fintech company in month t and district d. The period considered is the six months before and after May 2019 (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 17. 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

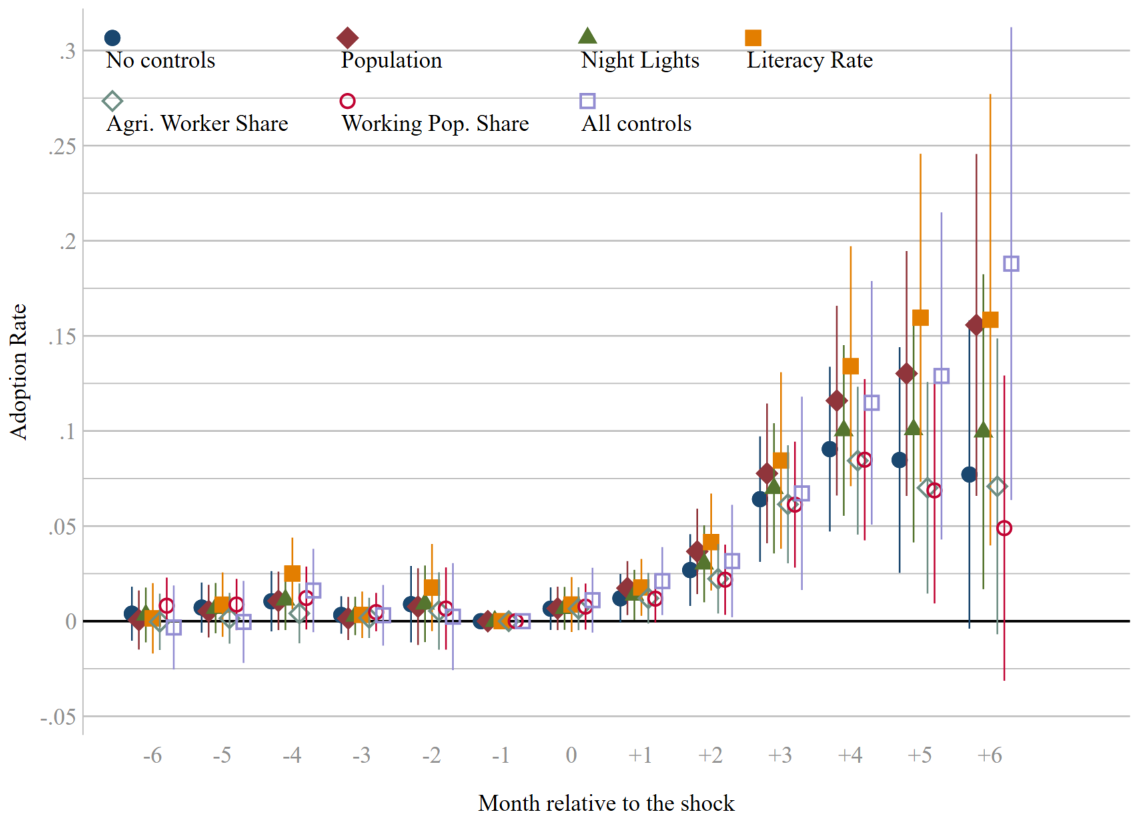

Notes: This figure reports a robustness of our main specification to the inclusion of controls. In particular, we reproduce the same Figure 3 with different level of controls. As in the main figure, the dependent variable is the number of stores that adopted our fintech company in month t and district 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

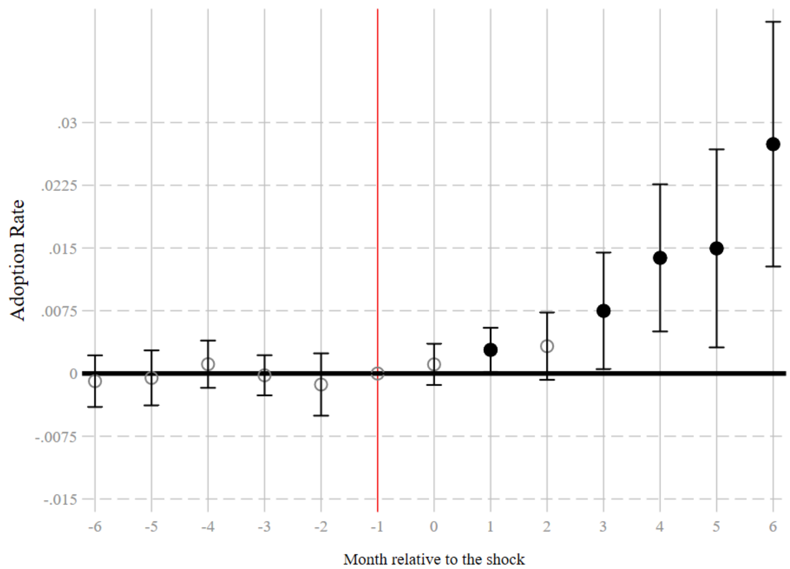

Notes: This figure provides a robustness test for the main dynamic specification. Everything is identical to the main figure (i.e., Figure 3), but for the treatment variable. In the main analysis, the treatment variable is the number of individuals between 15 and 29, scaled by the number of adults, defined as individuals between 15 and 74. This robustness figure instead uses as treatment a measures 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 17. 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

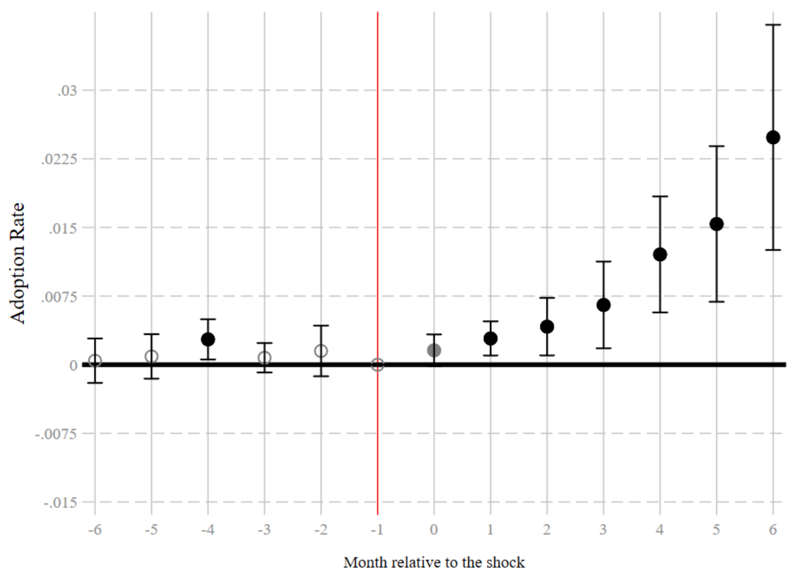

Notes: This figure provides a robustness test for the main dynamic specification. Everything is identical to the main figure (i.e., Figure 3), but for the treatment variable. In the main analysis, the treatment variable is the number of individuals between 15 and 29, scaled by the number of adults. This 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 17. 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))

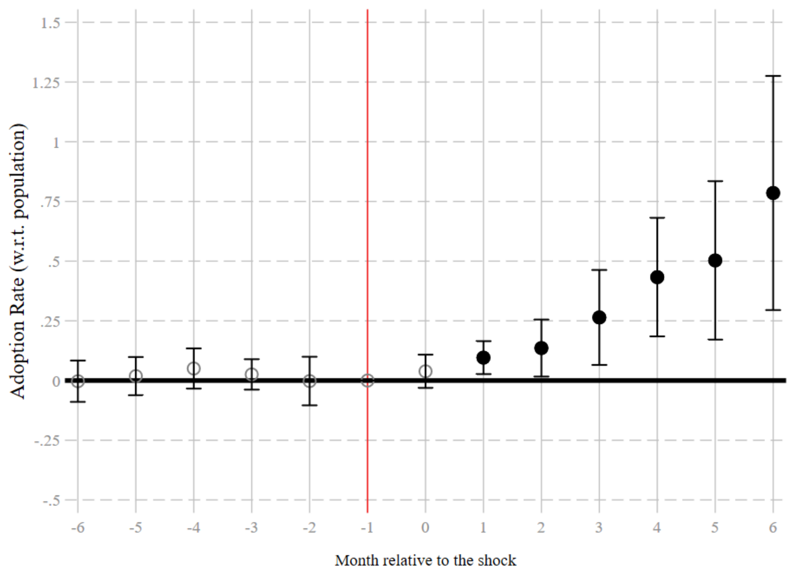

Notes: This figure provides a robustness test for the main dynamic specification. Everything is identical to the main figure (i.e., Figure 3), but for the way the outcome is constructed. In particular, the dependent variable is the number of stores that adopted our fintech company in month t and district 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 17. 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))

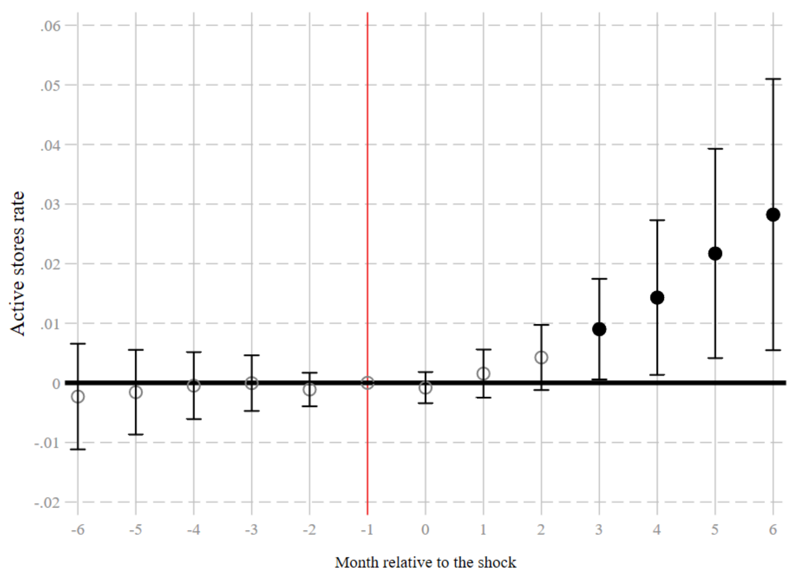

Notes: The figure plots the dynamic treatment effects of age structure on the total number of firms that are in our fintech platform. Apart from the outcome, everything is identical to the main figure (i.e., Figure 3). The dependent variable is the number of stores are in the platform in month t and 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 17. 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)

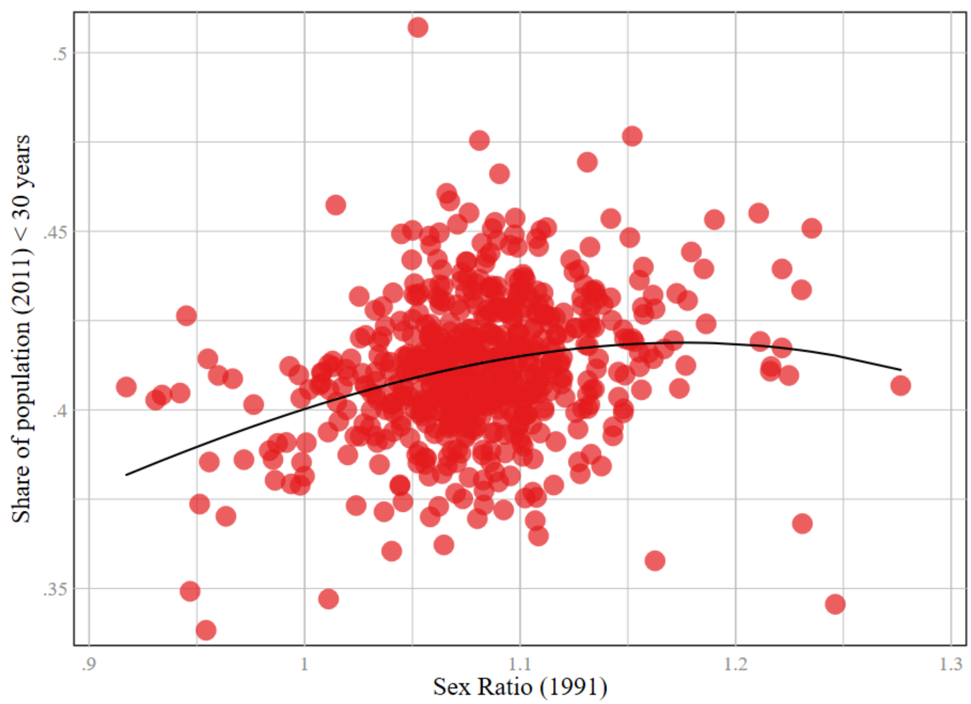

Notes: The figure documents the relationship between district sex ratios, defined as the number of males per female, in 1991 and the share of 2011 population in districts that is below 30 years. Each dot represents a district in the 2011 Census. The black line represents a quadratic polynomial fit.

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

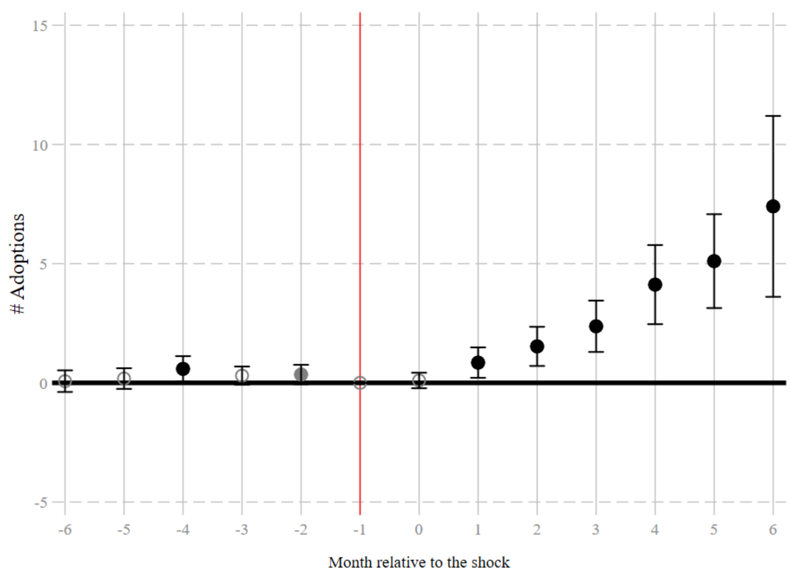

Notes: The figure provides a robustness to the main university result (Figure 5). In particular, we change the way the outcome is measure. The dependent variable is the (raw, without any transformation) the number of stores that adopted our fintech company at pincode-level in a month. The graphs report the coefficients γ_k from specification 18, 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

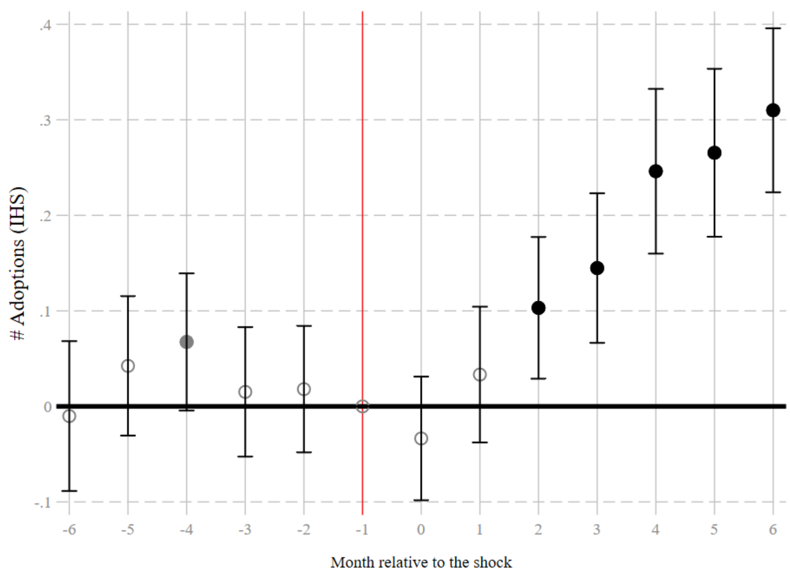

Notes: The figure provides a robustness to the main university result (Figure 5). In particular, everything is identical to the main analysis, with the exception that here we only use the sample of districts that have at least one university in their terriroty. The dependent variable is the (inverse 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 18, 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

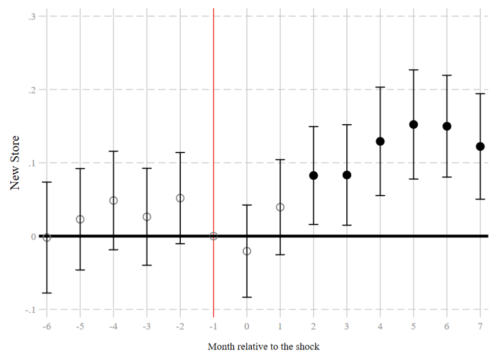

Panel (b): Student-consumer businesses, broader

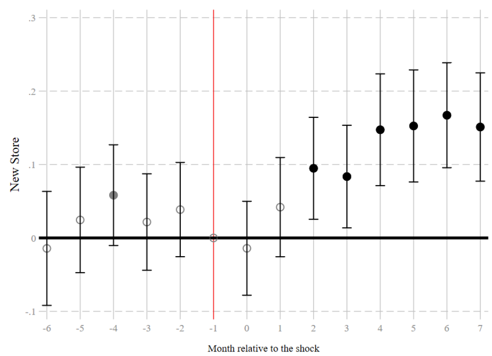

Panel (c): Placebo Merchants

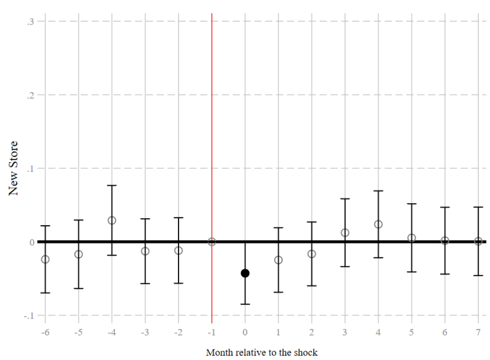

Notes: The figure provides a robustness to the main university result (Figure 5). Across the three panels, we replicate our main analyses using only a sub-sample of merchants. In panel (a), we identify businesses that are mostly depending on local demand by students (i.e., retailers, gas stations, restaurants, 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 18, 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

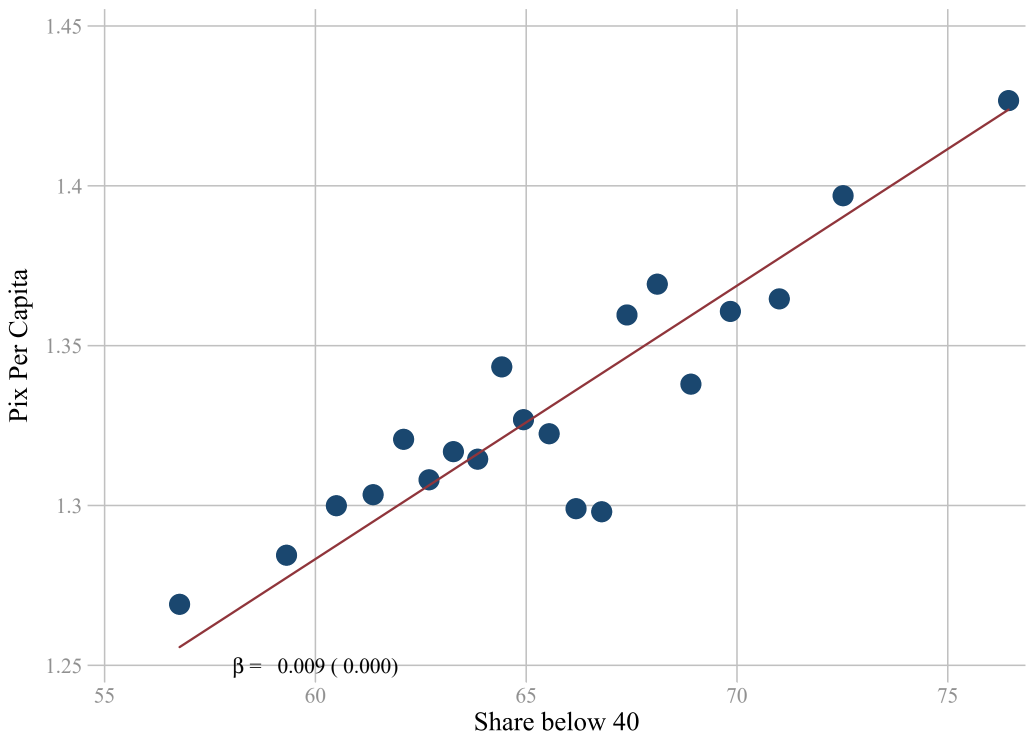

Notes: The figure examines the relationship between the use of Pix in Brazilian municipalities and the age structure of the municipalities, following a similar graphical representation presented before. The use of Pix is measured as the average Pix used (in value) per unit of population in Brazil. Pix 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***	0.048***
	(0.012)	(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(has university)$_p$× Post$_t$	0.065***	0.076***	0.000	0.171***
	(3.30)	(3.75)	(0.00)	(6.93)
Pincode FE	Y	Y	Y	Y
District × 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(has university)$_p$× Post$_t$	0.291***	0.332***	-0.004	2.494***
	(2.64)	(2.65)	(-0.24)	(4.42)
Pincode FE	Y	Y	Y	Y
District × 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.