Title
How Debit Cards Enable the Poor to Save More

Journal
The Journal of Finance, 76(4)

ISSN
0022-1082 1540-6261

Authors
BACHAS, PIERRE
GERTLER, PAUL
HIGGINS, SEAN
et al.

Publication Date
2021-05-07

DOI
10.1111/jofi.13021

The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXXVI, NO. 4 • AUGUST 2021

How Debit Cards Enable the Poor to Save More

PIERRE BACHAS, PAUL GERTLER, SEAN HIGGINS, and ENRIQUE SEIRA

ABSTRACT

We study an at-scale natural experiment in which debit cards were given to cash transfer recipients who already had a bank account. Using administrative account data and household surveys, we find that beneficiaries accumulated a savings stock equal to 2% of annual income after two years with the card. The increase in formal savings represents an increase in overall savings, financed by a reduction in current consumption. There are two mechanisms. First, debit cards reduce transaction costs of accessing money. Second, they reduce monitoring costs, which led beneficiaries to check their account balances frequently and build trust in the bank.

A REMARKABLY LARGE NUMBER OF households do not have sufficient savings to cope with relatively small shocks. For example, more than 40% of Americans

Pierre Bachas is at Development Research Group, World Bank. Paul Gertler is at Haas School of Business, UC Berkeley. Sean Higgins is at Kellogg School of Management, Northwestern University. Enrique Seira is at Department of Economics, ITAM. We are grateful to officials in Mexico's government bank Bansefi and the conditional cash transfer program Prospera (formerly Oportunidades) for sharing data and answering numerous questions. At Bansefi, we are indebted to Gabriel Cabañas, Miguel Ángel Lara, Oscar Moreno, Ramón Sanchez, and especially Benjamín Chacón and Ana Lilia Urquieta. At Prospera, we are indebted to Martha Cuevas, Armando Gerónimo, Rodolfo Sánchez, Carla Vázquez, and especially Rogelio Grados, Raúl Pérez, and José Solis. For comments that greatly improved the paper, we thank David Atkin, Alan Barreca, Richard Blundell, Chris Carroll, Carlos Chiapa, Shawn Cole, Natalie Cox, Pascaline Dupas, John Edwards, Gerardo Escaroz, Fred Finan, Jess Goldberg, Emilio Gutierrez, Jens Hainmueller, Anders Jensen, Anne Karing, Dean Karlan, Supreet Kaur, Leora Klapper, David Laibson, Ethan Ligon, John Loeser, Nora Lustig, Jeremy Magruder, Justin McCrary, Atif Mian, Ted Miguel, Doug Nelson, Christine Parlour, Jaime Ruiz-Tagle, Betty Sadoulet, Todd Schoellman, Jonathan Zinman, and numerous seminar participants. We are also grateful to Ignacio Camacho, Ernesto Castillo, Oscar Cuellar, Bernardo García, Austin Farmer, Joel Ferguson, and Isaac Meza for research assistance. Gertler and Seira gratefully acknowledge funding from the Consortium on Financial Systems and Poverty and the Institute for Money, Technology & Financial Inclusion. Higgins gratefully acknowledges funding from the Fulbright–García Robles Public Policy Initiative and National Science Foundation (Grant Number 1530800). The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent. The authors have read *The Journal of Finance*'s disclosure policy and have no conflicts of interest to disclose beyond the data and funding sources described above.

Correspondence: Sean Higgins, Kellogg School of Management, Northwestern University; e-mail: sean.higgins@kellogg.northwestern.edu.

DOI: 10.1111/jofi.13021

report that they "either could not pay or would have to borrow or sell something" to finance a $400 emergency (Federal Reserve (2017)). Some hypothesize that this is due to a lack of access to low-cost, convenient savings devices at formal financial institutions (Karlan, Ratan, and Zinman (2014)). When households do have access to financial institutions, they experience a number of well-documented causal impacts, including increased entrepreneurial investment, wealth accumulation, and ability to cope with shocks (Bruhn and Love (2014), Célérier and Matray (2019), Stein and Yannelis (2020)).

Nevertheless, take-up and active use of bank accounts "remain puzzlingly low" (Karlan et al. (2016, p. 2)), even when accounts are offered without fees (Dupas et al. (2018)). In fact, 40% of adults worldwide do not have a formal bank or mobile money account (Demirgüç-Kunt et al. (2018)). Similarly, cash transfer recipients paid through direct deposit into bank accounts generally withdraw the entire transfer amount in one lump sum each pay period (Muralidharan, Niehaus, and Sukhtankar (2016)).

We study a natural experiment in which debit cards tied to existing savings accounts were rolled out geographically over time to beneficiaries of the Mexican conditional cash transfer program Oportunidades. Debit cards alleviate two important barriers to using formal financial institutions. First, debit cards lower the indirect transaction costs of accessing money in an account by facilitating more convenient access via a network of ATMs. Second, debit cards reduce the indirect cost of checking balances, which enables individuals to verify that banks are not unexpectedly reducing balances. Through this monitoring channel, individuals build trust that money saved in a bank account will be available when wanted. In fact, a lack of trust in banks to not "steal" their savings—often through hidden and unexpected fees—is frequently listed as a primary reason the poor are hesitant to use banks (Dupas et al. (2016), FDIC (2016)). Among Oportunidades beneficiaries, "repeated balance checking is common, usually out of anxiety to confirm that their money is still there" (CGAP (2012, p. 20)).

The phased geographic rollout of debit cards to Oportunidades recipients provides plausibly exogenous variation in the timing of assignment of debit cards, which allows us to estimate the causal impact of having a debit card on saving in a difference-in-differences event study framework. Before the rollout, beneficiaries received their transfers through savings accounts without debit cards, and they rarely used their accounts to save—they typically withdrew the full transfer amount shortly after receiving it.[1]

Using high-frequency administrative data from nearly 350,000 beneficiary bank accounts in 357 bank branches nationwide over five years, we find that debit cards led to a large and significant increase in the active use of the accounts. The number of transactions (withdrawals) jumped immediately, while

[1] Prior to receiving cards, 13% of beneficiaries saved in the bank accounts. This is consistent with findings from other countries such as Brazil, Colombia, India, and South Africa, in which cash transfers are also paid through bank accounts and recipients generally withdraw the entire transfer amount in one lump-sum withdrawal each pay period (CGAP (2012), Muralidharan, Niehaus, and Sukhtankar (2016)).

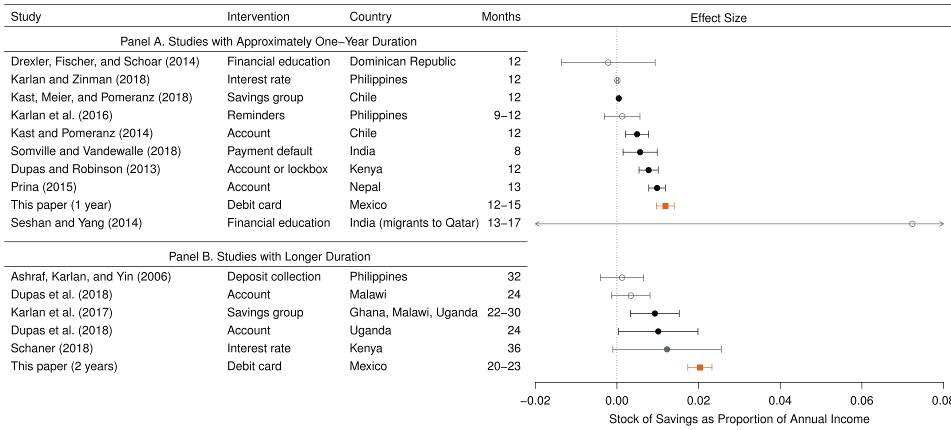

Figure 1. Comparison with other studies. This figure compares the results from our study after one and two years with a debit card (orange squares) to other studies of savings interventions, and shows that we find larger effects than most studies with comparable duration. Panel A focuses on studies with approximately one-year duration and Panel B focuses on studies with longer duration. The effect sizes are intent-to-treat effects of the intervention on the stock of savings, measured as a proportion of annual income. Internet Appendix Section I details the selection criteria to determine which studies to include and how we obtain their effects on the stock of savings as a proportion of annual income. Whiskers denote 95% confidence levels. Black (gray) circles indicate results that are signficant at the 5% (10%) level, while hollow circles indicate results that are statistically insignificant from zero. (Color figure can be viewed at wileyonlinelibrary.com)

the proportion of beneficiaries holding significant positive savings in their bank account increased more slowly, from 13% to 87% over a two-year period. After two years, beneficiaries with debit cards built up a stock of savings equal to 2% of annual income. This increase in savings—caused by an at-scale intervention that could be feasibly replicated with cash transfer beneficiaries in other countries—contrasts with the smaller and sometimes null effects on savings found by other interventions in the literature (Figure 1).

Using a rich household panel survey covering a subsample of the beneficiaries, we next test whether the increase we observe in formal savings is an increase in *overall* savings or a substitution from other forms of saving, both formal and informal. We find that after about one year with the card, there was no change in income and a significant reduction in consumption equal to about 4.9% of income. Because consumption and income are flows, and because the administrative bank account data show that the savings stock did not evolve linearly over time, we carefully compare this reduction in consumption of 4.9% of income to the change in the savings rate for beneficiaries from the same localities after they had the card for the same amount of time as in the survey. This change in the savings rate from the comparable administrative data is 4.6% of income.

These results suggest that total savings likely rose by a similar amount as we observe in the administrative bank account data (assuming that the flow of total savings is income minus consumption). More precisely, the administrative data suggest an increase in the savings rate of 4.6% of income, and the survey estimates show no change in income and a consumption reduction of 4.9% of income, both for beneficiaries from the same localities who had the card for approximately one year. Furthermore, the point estimates from the two data sources are nearly identical (within 0.2% of income, or less than 50 cents per month) and each lies within the 95% confidence interval of the other. As with the most household surveys, however, our survey estimates are noisy: the lower bound of the 95% confidence interval in the survey is 1.0% of income. Thus, while we can reject the hypothesis that the increase in formal savings was *purely* substitution from other forms of saving, we cannot rule out the possibility that *part* of the formal savings increase was substitution.

Debit cards could lead to increased total savings if saving informally is difficult and debit cards make saving formally more attractive. We find evidence consistent with it being difficult to save informally due to household members' easier access to the savings (consistent with evidence from lab-in-the-field experiments in Ashraf (2009) and Jakiela and Ozier (2016)). First, we find suggestive evidence of a larger proportional drop in spending on temptation goods compared to other goods, but no change to investment in education and health or assets. Second, we find suggestive evidence that after receiving a debit card, beneficiaries with low intrahousehold bargaining power at baseline increased savings more than beneficiaries with high intrahousehold bargaining power.

How do debit cards make saving formally more attractive? An obvious candidate is that debit cards decrease the transaction costs of accessing money, which makes saving in the account more attractive since savings can be easily accessed when needed. Indeed, debit cards reduced the indirect transaction costs of accessing the account: before receiving a card, account holders had to go to one of only 500 Bansefi branches nationwide to withdraw money, traveling a median road distance of 4.8 km.[2] After receiving the card, each beneficiary could withdraw her balance from *any* bank's ATM, that is, at any of the more than 27,000 ATMs in Mexico; she could also use the debit card to make purchases at point-of-sale (POS) terminals. The median road distance between a beneficiary's house and the closest ATM was 1.3 km (Bachas et al. (2018)). We find that the number of withdrawals made per month jumped by 36% immediately after receiving the card and stayed relatively flat thereafter. Many beneficiaries started making two or three withdrawals per transfer period, while almost all beneficiaries made a single withdrawal of the entire transfer prior to receiving a card. Furthermore, 16% of beneficiaries began accumulating savings immediately, likely due to the immediate reduction in transaction costs to access their money.

[2] This may explain their low initial use of the accounts to save. If clients were already saving in their accounts and the transaction costs provided a form of commitment device, as was the case for one of the households profiled by Morduch and Schneider (2017), it is possible that a reduction in transaction costs would reduce savings.

However, upon receiving a debit card, most beneficiaries did not begin saving immediately, but instead appear to have first used the card to monitor account balances and thereby build trust that their money was safe. Although a beneficiary could check her balance prior to receiving a card by going to a Bansefi branch and asking a bank teller, the debit card made balance checks much more convenient since they could be done at any bank's ATM. Thus, a reduction in transaction costs enabled trust-building. Once trust was established, beneficiaries took advantage of the reduced transaction costs of accessing money associated with debit cards and increased the amount of savings held in their bank accounts.[3]

Two main pieces of evidence support the mechanism of using the card to monitor balances and build trust. First, using the high-frequency administrative data on bank account transactions, we observe that upon receipt of the debit card, beneficiaries initially left small amounts of money in the account and used the card to check their account balances frequently, but reduced balance check frequency over time. We show that the reduction in balance checks over time was not driven by checking whether the transfer had arrived or whether there was enough money in the account before using the debit card to make a transaction at a POS terminal (furthermore, the Bansefi accounts did not charge overdraft fees). Second, in survey data from a subsample of the beneficiaries, those who had their debit cards for a short period of time reported significantly lower rates of trust in the bank than beneficiaries who had their debit cards longer. We also rule out a number of competing mechanisms including falling transaction costs over time and learning the banking technology, among others.

Our main contribution to the literature is to show that a nationwide, at-scale rollout of a low-cost financial technology caused a large and significant increase in the number of active account users in terms of both number of withdrawals and savings. The stock of savings accumulated after two years corresponds to 2% of annual income. This is larger than estimates from most other savings interventions—including offering commitment devices, no-fee accounts, higher interest rates, and financial education (Figure 1). Two other studies that also find a large effect on savings are Suri and Jack (2016), who study the impact of mobile money, and Callen et al. (2019), who study the impact of weekly home visits by a deposit collector equipped with a POS terminal. Like debit cards, these technologies both reduce transaction costs and enable clients to more easily monitor account balances (although these studies do not directly document the importance of these two channels).[4]

[3] The reduced indirect transaction costs of accessing money in the account and the ability to use the debit card for purchases may have also increased the benefit of saving formally. These factors may have increased beneficiaries' desire to learn whether the bank was trustworthy and the amount they decided to save once they trusted the bank.

[4] Mobile money clients can easily check account balances from their phones, and Callen et al.'s (2019) deposit collection included a receipt printed in real time with the deposit amount and new account balance after each weekly deposit—a feature that the bank viewed as crucial to establish trust in the deposit collectors. We are unable to include these studies in Figure 1 for reasons

Unlike debit cards, however, these technologies involve large implementation costs: given the existing ATM infrastructure in most countries, debit cards are very low-cost, while mobile money requires setting up an infrastructure of mobile money agents throughout the country and sending deposit collectors is labor-intensive.[5] Our paper also goes beyond these studies by showing that savings in the account were new savings financed by a reduction in consumption, and by providing evidence that both lower transaction costs and account monitoring explain the savings increase.

Our second contribution is to show that the savings effect came—at least in part—from an increase in total savings achieved by reducing current consumption, rather than a substitution from other forms of saving. Other studies testing whether an increase in formal savings represents an increase in total savings or substitution from informal savings do not typically have sufficient power to rule out full substitution, even when they find large point estimates on total savings (e.g., Ashraf et al. (2015), Kast, Meier, and Pomeranz (2018)).[6] While Somville and Vandewalle (2018, Figure 2) find that account holders may have reduced consumption as they increased savings over time, their consumption results are not statistically significant; Breza and Chandrasekhar (2019) similarly find noisy consumption results suggestive of a reduction in consumption to finance savings. In this paper, we definitively show that part (and based on the point estimates, possibly all) of the increase in formal savings was financed by a reduction in current consumption.

Our third contribution is to directly investigate two barriers to saving: indirect transaction costs and distrust. We find that some beneficiaries began saving immediately after receiving a debit card—likely due to the decreased transaction costs of accessing the account—while others began saving only after a delay. This delay was due at least in part to beneficiaries first monitoring the bank by checking account balances and increasing their trust in the bank over time. Prior studies explore the role of trust in stock market participation, in the use of checks instead of cash, and in mortgage refinancing in developed countries (Guiso, Sapienza, and Zingales (2004), Guiso, Sapienza, and Zingales (2008), Johnson, Meier, and Toubia (2019)), as well as the role

explained in Section I of the Internet Appendix. The Internet Appendix is available in the online version of the article on *The Journal of Finance* website.

[5] To our knowledge, Schaner (2017) is the only other paper that evaluates the impact of debit cards on savings. Her setting is very different from ours. First, recipients of the card (and the control group that received an account but no card) were a selected group—they had already expressed interest in opening an account at the partner bank. Second, in the rural Kenyan town in which her experiment was conducted, there was only one ATM located just outside one of the bank's branches. Thus, providing debit cards did not reduce travel costs to access money in the account or monitor account balances. Instead, withdrawing money at the ATM had a lower fee than interacting with a bank teller. She finds an increase in the number of transactions and the value deposited and withdrawn, but no change in savings.

[6] An exception is Callen et al. (2019), who find a statistically significant impact on total savings that is similar in magnitude to the impact on formal savings. Unlike our paper, they find no impact on consumption; instead, they find that an increase in labor supply in response to the savings intervention enabled the increase in savings.

of trust in borrowing and in the take-up of insurance products in developing countries (Karlan et al. (2009), Cole et al. (2013)). Few studies, however, rigorously explore the role of distrust as a constraint to saving or the role of financial technology in increasing trust (Karlan, Ratan, and Zinman (2014)).[7]

In summary, debit cards combined with ATMs or POS terminals are low-cost technologies that reduce the indirect transaction costs of both accessing funds in an account and checking balances to build trust in financial institutions. These technologies are simple, prevalent, and potentially scalable to hundreds of millions of households worldwide; they could be especially useful to enable the poor to save when combined with direct deposits (Blumenstock, Callen, and Ghani (2018)). In particular, government cash transfer programs could be a promising channel to increase financial inclusion, not only because of the sheer number of people that are served by cash transfers, but also because many governments and nongovernmental organizations are already embarking on digitizing their cash transfer payments through bank or mobile money accounts (e.g., Haushofer and Shapiro (2016), Muralidharan, Niehaus, and Sukhtankar (2016)). When these enabling conditions—direct deposits of sizable transfers and an expansive ATM network—are not present, the effect of debit cards on savings could be lower; our effect size is likely an upper bound for such contexts.

The rest of the paper is organized as follows. Section I describes the institutional context. Section II presents the data. Section III details the identification strategy and econometric specifications. Section IV measures the impact of debit cards on the savings and transactions of poor Mexican households. Section V shows that the increase in households' savings in their bank accounts corresponds to an increase in total savings. Section VI explores mechanisms behind the increase in savings, in particular that of building trust in the bank. Section VII concludes.

I. Institutional Context

We examine the rollout of debit cards to urban beneficiaries of Mexico's conditional cash transfer program Oportunidades, whose cash benefits were already being deposited directly into formal savings accounts without debit cards. Oportunidades is one of the largest and most well-known conditional cash transfer programs worldwide, with a history of rigorous impact evaluation (Parker and Todd (2017)). The program provides cash transfers every two months ("bimester") to poor families, conditional on sending their children to school and having preventive health checkups; due to program rules, the

[7] Previous studies on debit cards and mobile money focus on the effect of the lower transaction costs facilitated by these technologies to make purchases, access savings and remittances, and transfer money (Zinman (2009), Jack and Suri (2014), Schaner (2017)), but not their capacity to monitor and build trust in financial institutions. Two studies on trust and savings are Osili and Paulson (2014), who study the impact of past banking crises on immigrants' use of banks in the United States, and Mehrotra, Somville, and Vandewalle (2021), who promote interactions with bankers and find that account savings is strongly associated with trust in one's own banker.

payments are made directly to women in nearly all beneficiary households. It began in rural Mexico in 1997 under the name Progresa, and later expanded to urban areas as Oportunidades starting in 2002. By 2011, nearly one-fourth of Mexican households received benefits from Oportunidades, and in 2014, it was rebranded as Prospera.[8]

As it expanded to urban areas from 2002 to 2005, Oportunidades opened savings accounts in banks for beneficiaries in a portion of urban localities, and began depositing the transfers directly into those accounts. By 2005, beneficiary families in over half of Mexico's urban localities were receiving their transfer benefits directly deposited into savings accounts at Bansefi, a government bank created to increase savings and financial inclusion among underserved populations. The Bansefi savings accounts had no minimum balance requirement or monthly fees and paid essentially no interest.[9] No debit or ATM cards were associated with the accounts, so beneficiaries could only access their money at Bansefi bank branches. Because there were only about 500 Bansefi branches nationwide and many beneficiaries lived far from their nearest branch, accessing their accounts involved large transaction costs. The median urban household receiving its transfers in a Bansefi account was 4.8 km from the nearest Bansefi branch (Bachas et al. (2018)). Overall, the savings accounts were barely used prior to the introduction of debit cards: 89.9% of beneficiaries made one withdrawal each bimester, withdrawing 99.5% of the transfer on average (Table IA.I).

In 2009, the government began issuing Visa debit cards to beneficiaries who were receiving their benefits directly deposited into Bansefi savings accounts. The cards enabled account holders to withdraw cash and to check account balances at any bank's ATM, as well as make electronic payments at any store accepting Visa. Overdrafting was not permitted and there were no fees for attempting to overdraft: if the account had insufficient funds when attempting to make an ATM withdrawal or POS transaction, the transaction would not go through and an "insufficient funds" message would be displayed. Beneficiaries could make two free ATM withdrawals per bimester at any bank's ATM; additional ATM withdrawals were charged a fee that varied by bank. When Bansefi distributed the debit cards, they also provided beneficiaries with a training session on how and where to use the cards (Internet Appendix Section III). The training sessions did not vary over time and did not discuss savings, nor encourage recipients to save.

Our sample consists of urban beneficiaries who received their transfer benefits in bank accounts prior to the rollout of debit cards. As shown in Figure 2, Panel A, beginning in January 2009, debit cards tied to these existing bank

[8] The program has led to increases in school attendance and grade completion, improvements in childrens' health outcomes, modest increases in household consumption and caloric intake, and no change in labor market participation or labor market income (see Parker and Todd (2017) for a review). Note that these effects do not confound our study since everyone in our analysis is a beneficiary.

[9] Nominal interest rates were between 0.09% and 0.16% per year compared to an inflation rate of around 5% per year during our study period.

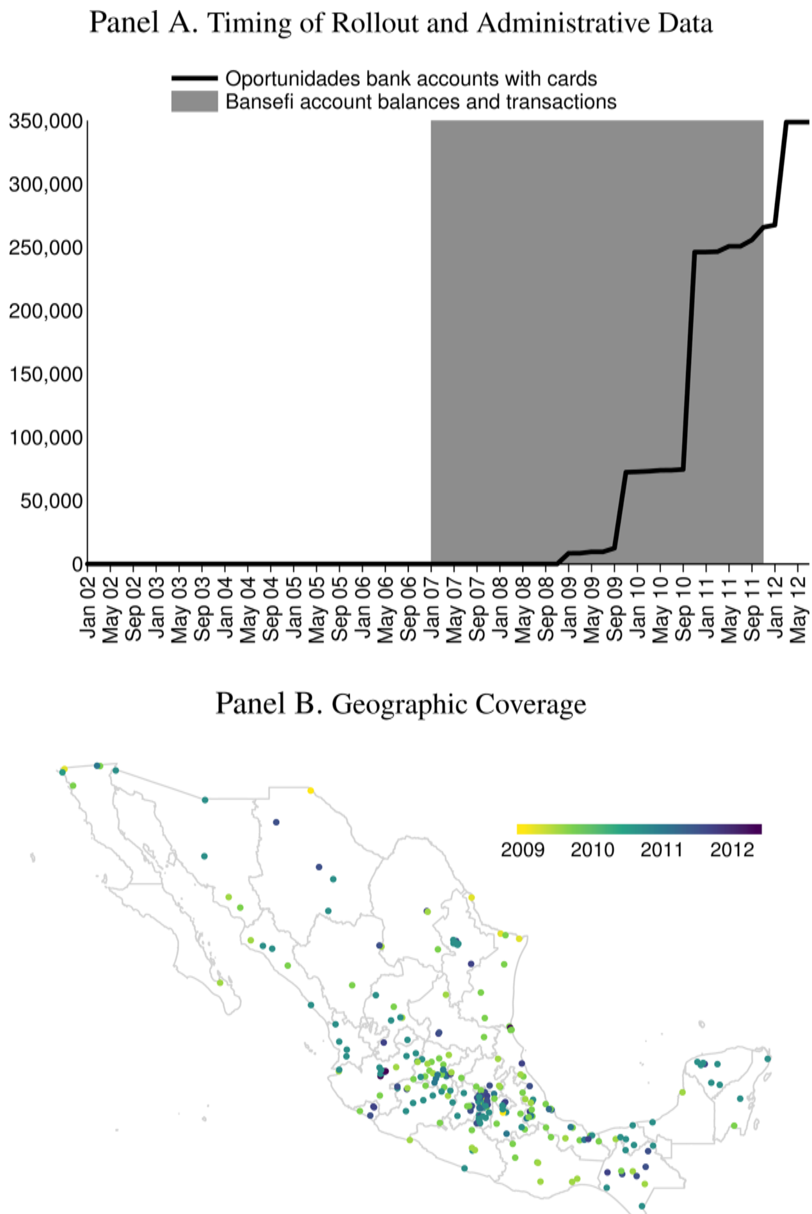

Figure 2. Debit card rollout over time and space. This figure shows the number of Oportunidades bank accounts with debit cards over time (using administrative data from Bansefi) and across space (using administrative data from Oportunidades). This was determined by the staggered rollout of debit cards, which generated variation over time and space in having a debit card tied to the bank account in which beneficiaries received their benefits. Panel A compares the timing of the rollout to the timing of the administrative bank account data and Panel B depicts the rollout across space. (Color figure can be viewed at wileyonlinelibrary.com)

accounts were rolled out to beneficiaries by locality. When Bansefi distributed cards in a particular locality, all beneficiaries in that locality received cards during the same payment period. By the end of 2009, about 75,000 beneficiaries had received debit cards tied to their preexisting savings accounts. Another 172,000 beneficiaries received cards by late 2010. By October 2011, the last month for which we have administrative data from Bansefi, a total of 256,000 beneficiaries had received debit cards tied to their preexisting savings accounts. Another 93,000 beneficiaries received cards between November 2011 and April 2012, shortly after the end date of our study period. We use this last group as a "pure" control group throughout the duration of our study, although as we describe in Section III, we take advantage of all the variation in exposure time generated by the staggered rollout of cards at the locality level over time. The map in Figure 2, Panel B, shows that the card expansion had substantial national geographic breadth throughout the rollout.

While the timing of the card rollout was not explicitly randomized across localities, we show in Section III that the timing was uncorrelated with observed locality-level characteristics. In conversations we conducted with Oportunidades officials, they explained that they did not target localities with particular attributes because they wanted to test their administrative procedures for the rollout—such as how easy it would be to distribute cards—on a quasi-representative sample. In addition, we show that the variables we use from the transaction-level data, as well as other variables such as wages, prices, and financial infrastructure, exhibit parallel pre-trends. Importantly, when a locality was treated, all beneficiaries in the locality received cards that period. Furthermore, although the total number of beneficiaries increased slightly over time at the national level, we show that the rollout was not accompanied by a differential change in the number of beneficiaries in a locality (Section III).

II. Data Sources

We use four main sources of data. The first is administrative data on account balances and transactions from Bansefi on the universe of beneficiaries who received benefits in a savings account and were then given a debit card. We also use three surveys of Oportunidades beneficiaries. Table I displays the number of beneficiaries, time periods, main variables, and variation we exploit for each of these data sources.

A. Administrative Data

To examine the effect of debit cards on savings and account use, we exploit account-level data from Bansefi on the average monthly balance and all transactions for the universe of accounts that received transfers in a savings account prior to receiving a debit card. These data consist of 348,802 accounts at 357 Bansefi branches over almost five years, from January 2007 to October 2011. They include the monthly average account balance; the date, amount, and type of each transaction made in the account (including Oportunidades

Table I

Summary of Data Sources and Identification

This table presents details on the four main data sources used in the paper.

Data Source	# Benef.	Period	Main Variables	Variation Used
(1) Administrative bank account data from Bansefi	348,802	Continuous panel: Jan 2007 to Oct 2011	Balances, transactions, balance checks	Generalized difference-in-differences (event study with control) using phased geographic rollout
(2) *Household Panel Survey* from Oportunidades (ENCELURB)	2,868	Panel (four waves): 2002, 2003, 2004, and November 2009 to February 2010	Consumption, income, assets	Difference-in-differences: received card in 2009 versus received card later
(3) *Trust Survey* from Oportunidades (ENCASDU)	1,694	Cross-section: October to November 2010	Self-reported reasons for not saving: for example, lack of trust, lack of knowledge	Tenure with card below/above median time in survey (median = 14 months)
(4) *Payment Methods Survey* from Oportunidades	1,617	Cross-section: June 2012	Self-reported number of balance checks, knowledge of technology	Tenure with card below/above median time in survey (median = 12 months)

transfers); the date the account was opened; and the month the card was given to the account holder. Figure 2, Panel A, shows the timing of the administrative data and the rollout of debit cards.

Table IA.I shows summary statistics from this data set. Using data from the first bimester of 2008 (before any debit cards were disbursed to beneficiaries), the accounts in our sample made 0.01 client deposits and 1.1 withdrawals per bimester on average. The average amount withdrawn was 99.5% of the Oportunidades transfer, indicating very low use of the account for saving prior to receiving the card. End-of-period balances were 124 pesos or about US$11 on average; the distribution of end-of-period balances is skewed: the 25$^{\text{th}}$ percentile is just 2 pesos (US$0.20) and the median is 42 pesos (US$4). The average amount transferred by Oportunidades in the first bimester of 2008 was 1,540 pesos, or about US$144; using survey data, we find that Oportunidades income represented about one-fourth of beneficiaries' total income on average. The average account had already been open for 3.5 years by January 2008, so beneficiaries in our study had substantial experience with a savings account prior to receiving the debit card.

B. Survey Data

Since its inception in 1997, Oportunidades has a long history of collecting high-quality surveys from their beneficiaries, and these surveys have been used extensively by researchers (Parker and Todd (2017)). We use three distinct Oportunidades household-level surveys, described below. Figure IA.1 shows when survey respondents received cards in each of these surveys, relative to the timing of the survey. In all surveys, the sample we use for estimation consists of households who received cards at some point during the rollout; these households correspond to a subset of the accounts in the administrative data described in Section II.A. Note that we cannot merge the survey data to the administrative account data.

B.1. Household Panel Survey (ENCELURB)

The most comprehensive survey we use is the *Encuesta de las Características de los Hogares Urbanos* (*ENCELURB*), a household panel survey with comprehensive modules on consumption, income, and assets. The survey includes three pretreatment waves in 2002, 2003, and 2004, and one posttreatment wave conducted between November 2009 and February 2010. The surveys were originally collected for the evaluation of the program in urban areas. Localities that switched to debit cards in early 2009 were oversampled in the fourth wave (which did not return to all localities from the original sample for budgetary reasons). As a result, most of the treatment group in this survey—beneficiaries who received cards prior to the fourth wave of the survey—had the card for close to one year when surveyed. We exclude the small group of beneficiary households in this survey that received cards in late 2009, shortly before the posttreatment survey wave, for cleaner comparisons with the administrative data results.[10] We merge the survey with administrative data from Oportunidades on the debit card expansion (at the locality level) to study the effect of the card on consumption and saving in a difference-in-differences model.

B.2. Trust Survey (ENCASDU)

The *Encuesta de Características Sociodemográficas de los Hogares Urbanos* (*ENCASDU*), conducted in 2010, is a stratified random sample of 9,931 Oportunidades beneficiaries. We refer to this survey as the *Trust Survey* since it gives us a measure of trust in the bank. We restrict our analysis to beneficiaries who had already received debit cards by the time of the survey, since the module with questions we use about reasons for not saving was only asked of those who had already received debit cards. This leaves us with a sample of 1,694 households, with a median exposure to the card of 14 months.

[10] Because only 74 of the 2,942 households in this survey living in urban localities included in the rollout were in localities treated in late 2009, our results are virtually unchanged if we do not drop these households.

The survey asked, "Do you leave part of the monetary support from Oportunidades in your bank account?" If the response was no, the respondent was then asked the open-ended question, "Why don't you keep part of the monetary support from Oportunidades in your Bansefi savings account?" Lack of trust is captured by responses such as "because if I do not take out all of the money I can lose what remains in the bank," "because I don't feel that the money is safe in the bank," "distrust," and "because I don't have much trust in leaving it."[11]

B.3. Payment Methods Survey

The *Encuesta de Medios de Pago* (*Payment Methods Survey*) is a cross-sectional survey of a stratified random sample of 5,388 beneficiaries, conducted in 2012. This survey was fielded to measure operational details of the payment method. In particular, it asked about use of the debit cards and beneficiaries' experiences using ATMs. We use this survey to obtain information on the self-reported number of balance checks and withdrawals with the card, whether beneficiaries got help using an ATM, and whether beneficiaries knew their card's PIN. We restrict the analysis to the 1,617 surveyed beneficiaries in the sampled urban localities that received cards prior to the survey; median exposure time to the card is 12 months.

C. Auxiliary Data

We use auxiliary data for four purposes: to identify (in survey data) when beneficiaries in each locality received cards as part of the rollout, to test for balanced pretrends across a broad range of observables, to test whether the timing of the rollout is correlated with locality-level characteristics, and to test for a supply-side response by banks to the debit card rollout.

C.1. Auxiliary Administrative Data

We use administrative data from Oportunidades on the number of beneficiaries and payment method by locality $\times$ bimester beginning in 2007. These data allow us to identify the timing of the rollout of debit cards (which is not necessary for results using the Bansefi administrative data—where we directly observe when each account received a debit card—but that we use for results using survey data). We also use data from Mexico's National Banking and Securities Commission (CNBV), which include a number of financial indicators at the bank $\times$ municipality $\times$ quarter level beginning in the fourth quarter of 2008. From these data, we use the number of bank branches, ATMs, debit cards, and credit cards to test for balanced pretrends and to test for a supply-side response by banks. We use administrative data from Mexico's

[11] We also use this question to define alternative reasons for not saving, including lack of knowledge (e.g., "they didn't explain the process for saving") and fear of ineligibility (e.g., "because if I save in that account they can remove me from the Oportunidades program").

Central Bank on the universe of POS terminal adoptions and cancelations since 2006 (Higgins (2019)) to test for balanced pretrends in financial technology adoption on the supply side of the market. Finally, we use local elections data that we digitized to test whether the timing of the rollout was determined by political considerations (specifically, the party in power at the local level).

C.2. Auxiliary Survey Data

We use data from three surveys to test for balanced levels and pretrends in the economic performance of the localities. First, we use locality-level indicators derived from Mexico's 2005 Population Census. The indicators we use are the ones used by Mexico's National Council of Social Development and Policy Evaluation (CONEVAL) to measure locality-level development gaps. Second, we use microdata on wages from Mexico's quarterly labor force survey, the *Encuesta Nacional de Ocupación y Empleo* (*ENOE*). This data set includes wages for 20 million individual $\times$ quarter observations over 2005 to 2016. Third, we use price quotes from the microdata used to construct Mexico's consumer price index. These data include over 4 million price quotes from 2002 to 2014 at the product $\times$ store $\times$ month level for food, beverages, alcohol, and tobacco.

III. Empirical Strategy and Identification

We exploit variation generated by the staggered rollout of debit cards to different localities by Oportunidades. In this section, we show that conditional on being included in the rollout, the timing of when a locality received treatment is not correlated with levels or trends in observables from a number of data sets. These data sets include microdata on wages and food prices, locality-level data on financial infrastructure and poverty, local elections data, and transaction-level data from beneficiaries' bank accounts.

Our empirical strategy depends on the data being used, but the underlying variation we use always stems from the plausibly exogenous rollout of debit cards over time. When the data have a panel dimension (that is, the administrative data and the *Household Panel Survey*), we estimate a difference-in-differences specification. When we only have a cross-section of account holders (that is, the *Trust Survey* and *Payment Methods Survey*), we exploit variation in the length of time beneficiaries had been exposed to the card. In this section, we present the main empirical models we use and verify the plausibility of the identification assumptions needed for a causal interpretation.

A. Generalized Difference-in-Differences (Event Study)

The large sample over a long period of time in the administrative data allows us to estimate a generalized difference-in-differences specification where the treatment effect is allowed to vary dynamically over time and is measured in "event time" relative to each beneficiary's treatment period. In other words, we use an event-study specification with a pure control group throughout the study period, where the pure control group comprises those who were treated after October 2011, the last period for which we have data. Specifically, we estimate

$$y_{it} = \lambda_i + \delta_t + \sum_{k=a}^{b} \phi_k D_{it}^k + \varepsilon_{it}, \tag{1}$$

where y_{it} is the outcome of interest, i and t index account and period, respectively, λ_i are account-level fixed effects, δ_t are calendar-time (as opposed to event-time) fixed effects, D_{it}^k is a dummy variable indicating that account i has had a debit card for exactly k periods at time t, and $a < 0 < b$ are periods relative to the switch to debit cards. We measure effects relative to the period before getting the card, so we omit the dummy for $k = -1$. For those in the control group who received cards after our study period ends, $D_{it}^k = 0$ for all k.[12] We use this specification to study withdrawals and savings in the account. We average time over four-month periods since payments were sometimes shifted to the end of the previous bimester.[13] We estimate cluster-robust standard errors, clustering ε_{it} by locality.

As in any difference-in-differences model, to interpret each ϕ_k as the causal effect of having the card for k periods, we need to invoke a parallel trends assumption: in the absence of the card, early and late card recipients would have had the same changes in account use and savings behavior. While this is untestable, we test for parallel preintervention trends by showing that $\phi_k = 0$ for all $k < 0$. We perform these tests not only for the outcomes that we use with specification (1), which come from the Bansefi savings and transactions data, but also for a number of other outcomes from numerous data sources.

Figure 3, Panels A and B, shows that the timing of when different localities received cards as part of the rollout is not correlated with pretrends in wages, food prices, or financial technology (POS terminals, bank branches, ATMs, or debit and credit cards). Furthermore, Panel C shows that it is not correlated with beneficiary savings or with the number of withdrawals made from their accounts. Using less granular annual data, we also test and rule out the possibility that the rollout is correlated with pretrends in the number of program beneficiaries or with whether the party in power at the municipal level corresponded to the party in power at the national level (Figure IA.2). In addition to demonstrating parallel pretrends, Figure IA.2 shows that there was no differential change in the number of program beneficiaries or local politics as a result of the debit card rollout.

[12] Since we have a control group that did not receive cards until after the study period ends, we can pin down the calendar-time fixed effects without facing the underidentification problems described in Borusyak and Jaravel (2016). We set a and b as the largest number of periods before or after receiving the card that are possible in our data, but only graph the coefficients representing three years before receiving the card and two years after (following Borusyak and Jaravel (2016)).

[13] This could lead to an artificially large end-of-bimester balance if the recipient had not yet withdrawn her transfer. Payment shifting happened for various reasons, including local, state, and federal elections, as a law prohibited Oportunidades from distributing cash transfers during election months.

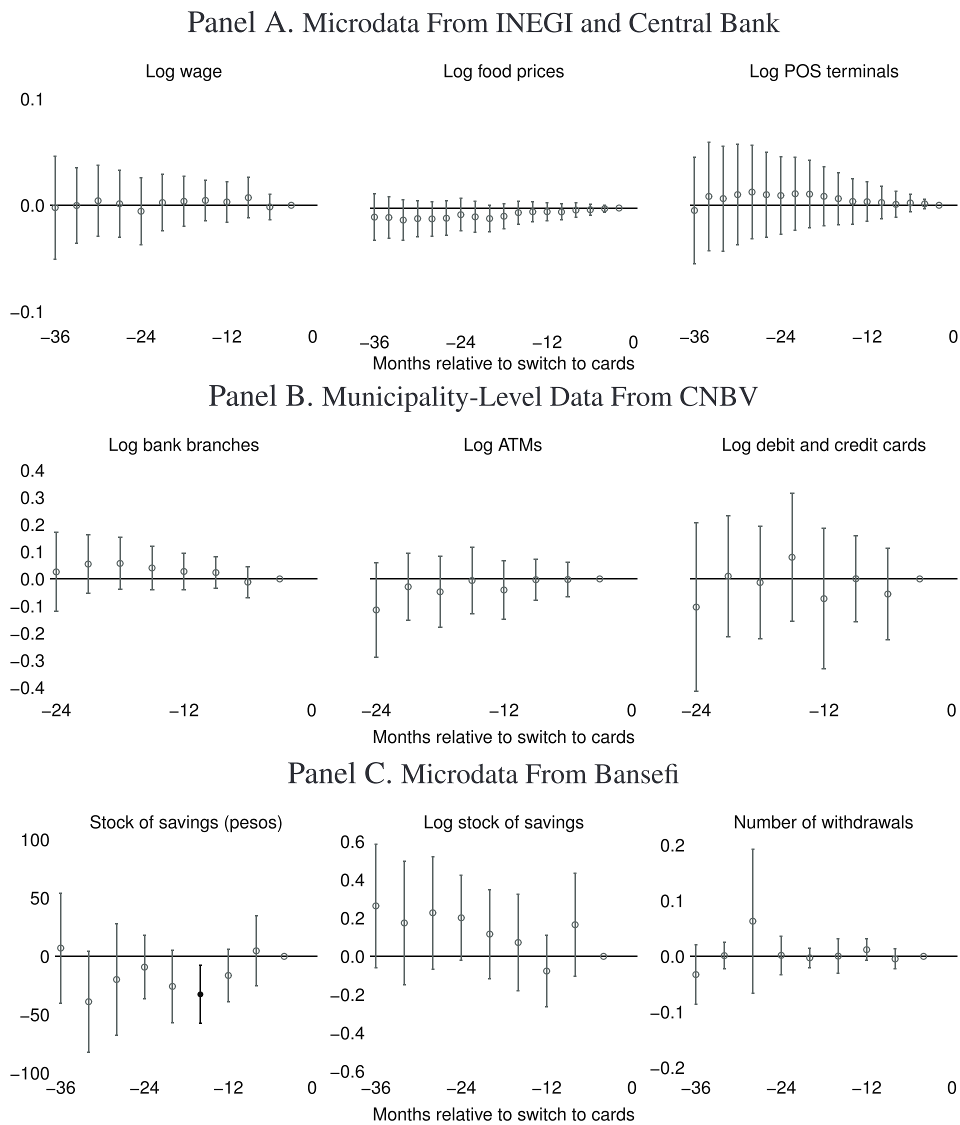

Figure 3. Parallel pretrends. This figure shows parallel pretrends in variables from (a) microdata on wages from INEGI's ENOE labor force survey, food prices from 4 million product × store × month price quotes from INEGI, and POS terminal adoptions from Mexico's Central Bank, (b) municipality-level data on financial variables from CNBV, and (c) microdata from beneficiaries' bank accounts. Point estimates are ϕ_k for $k < 0$ from specification (1), where $k = -1$ is the omitted period. In the wage regression, i in specification (1) is a worker; in the food price regression, i is a product × store; and in the POS terminal regressions, the data are aggregated to the postal code level and i is a postal code. In Panel B, the data are aggregated to the municipality level and i is a municipality. In Panel C, i is a Bansefi account. The frequency of ϕ_k coefficients depends on the frequency of each data set. Panel B includes 24 rather than 36 months of pretrends because the CNBV data begin in the last quarter of 2008, so the sample of localities with more than 24 months pretreatment is very small. Standard errors are clustered at the locality level in Panels A and C, and at the municipality level—since the data are only available by municipality, which is slightly larger than locality—in Panel B. Black circles indicate results that are significant at the 5% level, while hollow circles indicate results that statistically insignificant from zero.

As an additional test of whether the timing of the rollout is correlated with levels or trends in locality-level observables, we follow Galiani, Gertler, and Schargrodsky (2005) and use a discrete-time hazard model. This is equivalent to testing whether in a given period t, the probability of being treated at t conditional on not being treated yet at $t-1$ is correlated with observables. We combine several data sets to include measures of the pretreatment levels and trends of financial infrastructure, politics, and the locality-level variables used by CONEVAL—the independent government agency that produces Mexico's official poverty estimates.[14] We reject the possibility that the timing of the rollout is correlated with observables among localities included in the rollout: of the 22 variables included in the model, the coefficient on one variable is statistically significant at the 5% level (as expected by chance) and the remaining coefficients are statistically insignificant (Table II).

B. Difference-in-Differences with Survey Data

With the *Household Panel Survey* data, we estimate a standard difference-in-differences model since we observe just one time period after treatment. We estimate

$$y_{it} = \lambda_i + \delta_t + \gamma D_{j(i)t} + \nu_{it}, \tag{2}$$

where y_{it} is consumption, income, or the stock of assets for household i at time t. Time-invariant differences in household observables and unobservables are captured by the household fixed effects λ_i, common time shocks are captured by the time fixed effects δ_t, and $D_{j(i)t} = 1$ if locality j in which beneficiary household i lived prior to treatment received debit cards by time t. We use the locality of residence prior to treatment to avoid confounding migration effects, and we estimate cluster-robust standard errors clustered by locality.[15]

The identifying assumption is again parallel trends. In addition to the evidence that the rollout of cards is not correlated with levels or trends of variables from several data sets in Section III.A, we verify parallel pretreatment trends in the *Household Panel Survey* data by estimating

$$y_{it} = \lambda_i + \delta_t + \sum_k \omega_k T_{j(i)} \times \mathbb{I}(k = t) + \eta_{it}, \tag{3}$$

[14] We include trends for the variables for which it is possible, that is, those from data sets with at least an annual frequency (ruling out the 2005 Census) with data beginning prior to 2008. The data on Bansefi branches and ATMs begin in the last quarter of 2008, and thus, we only include levels of those variables.

[15] In our data, very few households migrated. In theory, this could be a result of migrating households attriting from the survey. Nevertheless, we confirm using other data that migration in these localities was low. Using data from a panel of 12 million voter registrations (a 15% random sample from the universe of 80 million voter registrations in Mexico), we check the proportion of residents from the same localities as those in the *Household Panel Survey* who migrated over a three-year period and find that only 4.5% of residents migrated to another locality.

Table II

Summary Statistics and Discrete-Time Hazard of Locality Characteristics

Columns (1) and (2) show summary statistics of locality-level financial infrastructure, trends in financial infrastructure, and other locality characteristics. Columns (3) and (4) test whether these characteristics predict the timing of when localities received debit cards as part of the debit card rollout, using a discrete-time hazard model (respectively, using linear probability and a proportional hazard using a complementary log-log regression). Both models also include a fifth-order polynomial in time as in Galiani, Gertler, and Schargrodsky (2005); time is measured in two-month periods. The dependent variable in the discrete-time hazard model is a dummy variable indicating if locality j has been treated at time t. A locality treated in period t drops out of the sample in period $t+1$ since it is a hazard model. All variables are measured prior to the debit card rollout. The financial variables are measured in the last quarter of 2008; prerollout trends (variables with a Δ) compare the last quarter of 2006 to the last quarter of 2008. The number of point-of-sale (POS) terminals come from Mexico's Central Bank, the number of checking accounts, commercial bank branches, and commercial bank ATMs comes from CNBV, and the number of Bansefi bank branches comes from Bansefi. We do not include trends in Bansefi bank branches or commercial bank ATMs because these variables are first available in 2008. The nonfinancial locality characteristics include all characteristics used to measure locality-level development in Mexico, and come from locality-level data from the 2005 Population Census published by INEGI. $N = 240$ localities in the debit card rollout, and 1,851 locality $\times$ two-month-period observations in columns (3) and (4).

			Discrete-Time Hazard	
Variable	Mean (1)	*SD* (2)	Linear Probability (3)	Proportional Hazard (4)
Log point-of-sale terminals	4.47	2.11	0.0002	0.0043
			(0.0095)	(0.0842)
Δ Log point-of-sale terminals	0.81	0.38	−0.0260	−0.2360
			(0.0185)	(0.1601)
Log bank accounts	9.27	3.27	0.0061	0.0537
			(0.0052)	(0.0435)
Δ Log bank accounts	1.78	3.61	0.0049	0.0495
			(0.0065)	(0.0558)
Log commercial bank branches	2.58	1.42	−0.0225	−0.2160
			(0.0187)	(0.1508)
Δ Log commercial bank branches	0.61	0.95	−0.0215	−0.2267
			(0.0240)	(0.2178)
Log Bansefi bank branches	0.58	0.41	0.0033	0.0420
			(0.0241)	(0.2001)
Log commercial bank ATMs	3.15	1.74	0.0130	0.1203
			(0.0103)	(0.0997)
Log population	11.26	1.24	0.0117	0.1072
			(0.0159)	(0.1317)
Mayor = PAN	19.58	39.77	−0.0003	−0.0027
			(0.0003)	(0.0023)
Δ Mayor = PAN	−12.08	57.67	0.0002	0.0021
			(0.0002)	(0.0016)
% illiterate (age 15+)	6.14	3.69	0.0004	0.0049
			(0.0048)	(0.0417)
% not attending school (age 6−14)	4.15	1.65	0.0003	0.0063
			(0.0094)	(0.0848)

(*Continued*)

Table II—*Continued*

Variable	Mean (1)	*SD* (2)	Discrete Time Hazard: Linear Probability (3)	Discrete Time Hazard: Proportional Hazard (4)
% without primary education (age 15+)	40.98	9.59	0.0018	0.0145
			(0.0019)	(0.0169)
% without health insurance	45.68	16.15	−0.0011	−0.0099
			(0.0008)	(0.0066)
% with dirt floor	5.28	4.83	0.0051**	0.0513**
			(0.0024)	(0.0209)
% without toilet	5.89	3.60	−0.0063	−0.0526
			(0.0040)	(0.0335)
% without water	6.45	9.12	−0.0007	−0.0058
			(0.0010)	(0.0094)
% without plumbing	3.94	6.39	0.0021	0.0180
			(0.0015)	(0.0122)
% without electricity	4.29	2.24	0.0052	0.0430
			(0.0048)	(0.0394)
% without washing machine	33.64	14.33	−0.0006	−0.0071
			(0.0010)	(0.0098)
% without refrigerator	16.80	9.73	0.0010	0.0068
			(0.0017)	(0.0153)

where k indexes survey round ($k = 2002$ is the reference period and is therefore omitted), $T_{j(i)} = 1$ if locality j in which beneficiary i lives is a locality that received cards before the posttreatment survey wave, and $\mathbb{I}(k = t)$ are time dummies. Thus, the ω_k coefficients for $k < 2009$ estimate placebo difference-in-differences effects for the pretreatment years. For each variable, we fail to reject the null of parallel trends using an F-test of $\omega_k = 0$ for all $k < 2009$ (Table III, Panel B, column (4)).

C. Cross-Section Exploiting Variation in Time with Card

The *Trust Survey* and *Payment Methods Survey* are cross-sections of beneficiaries with cards (thus there is no pure control group), and each survey has fewer than 2,000 observations. This poses constraints: we have to rely on exposure time to the card as the identifying variation, and to economize on power, we split the beneficiaries into two equal-sized groups based on how long they had the card. Note that the variation in time with the card is still determined by the plausibly exogenous rollout of cards by the government.

We regress the outcome variable—such as self-reported reasons for not saving—on a dummy for whether beneficiary i's time with the card is above the median,

$$y_i = \alpha + \gamma \mathbb{I}(\text{Card} \geq \text{median time})_i + u_i, \tag{4}$$

where u_i is clustered at the locality level.

Table III
Balance and Parallel Trends in Survey Data

This table tests for balance between those with above- versus below-median time with a debit card in the two cross-sectional surveys, and for parallel trends in the panel survey. Panel A shows results from specification (4): column (1) shows the mean for those with below-median time with a debit card (α), column (2) shows the difference in means for those with above-median time with a card (relative to those with below-median time with a card; γ), and column (3) reports p-values for a test of $\gamma = 0$. The *Payment Methods Survey* included fewer sociodemographic questions, which is why some rows are blank in the columns corresponding to that survey. $N = 1{,}694$ beneficiary households for the *Trust Survey* and 1,617 for the *Payment Methods Survey*. Panel B shows the control mean and a parallel trends test for each of the outcome variables used in the *Household Panel Survey*. The parallel trends test is from specification (3). The "Placebo diff-in-diff" columns show ω_{2003} and ω_{2004} ($k = 2002$ is the omitted reference period), while the "Parallel p-value" column is from an F-test of $\omega_{2003} = \omega_{2004} = 0$. $N = 7{,}754$ household-period observations from 2,200 households in the *Household Panel Survey* as in column (4) of Table V. Asymptotic standard errors clustered at the locality level are included in parentheses. Asymptotic cluster-robust p-values are included without brackets and randomization inference p-values based on 2,000 draws are in square brackets.

Panel A. Cross-Sectional Data

	Trust Survey			*Payment Methods Survey*		
	α (Mean for Card < Median Time) (1)	γ (Difference Card ≥ Median Time) (2)	p-Value of Difference (3)	α (Mean for card < Median Time) (4)	γ (Difference Card ≥ Median Time) (5)	p-Value of Difference (6)
# Household members	5.44	−0.26	0.11	4.74	0.04	0.77
	(0.12)	(0.15)	[0.16]	(0.13)	(0.13)	[0.78]
Age	45.53	−1.04	0.20	39.03	1.18	0.15
	(0.83)	(0.78)	[0.30]	(0.86)	(0.80)	[0.14]
Male	0.01	0.00	0.95	0.02	0.00	0.87
	(0.00)	(0.01)	[0.81]	(0.01)	(0.01)	[0.88]
Married	0.72	−0.03	0.33	0.72	0.01	0.87
	(0.02)	(0.03)	[0.42]	(0.03)	(0.03)	[0.87]
Education level	5.84	0.33	0.19	6.04	−0.03	0.91
	(0.16)	(0.24)	[0.27]	(0.28)	(0.29)	[0.91]
# Children	2.22	−0.03	0.74			
	(0.09)	(0.10)	[0.76]			
Occupants per room	3.48	0.03	0.80			
	(0.08)	(0.11)	[0.82]			
Health insurance	0.63	−0.05	0.16			
	(0.03)	(0.03)	[0.26]			
Asset index	0.01	0.05	0.53			
	(0.08)	(0.08)	[0.58]			
Income	3,443.60	−218.20	0.16			
	(136.42)	(149.03)	[0.23]			

(*Continued*)

Table III—*Continued*

Panel B. Panel Data				
	Household Panel Survey			
	Control	ω_k (Placebo diff-in-diff)		Parallel
	Mean (1)	2003 (2)	2004 (3)	p-Value (4)
Consumption	2,731.20	2.04	9.37	0.95
	(82.83)	(81.01)	(85.17)	[0.97]
Income	3,148.28	265.96	275.18	0.47
	(89.06)	(219.09)	(224.06)	[0.54]
Asset index	0.47	−0.03	−0.02	0.67
	(0.10)	(0.05)	(0.05)	[0.71]

This specification requires orthogonality between the error term u_i and timing of card receipt for a causal interpretation of γ—a stronger identification assumption than parallel trends.[16] We thus conduct balance tests using specification (4) with characteristics that should not be affected by debit card receipt as the dependent variable, such as number of household members, age, gender, marital status, education level, assets, and income. Table III, Panel A, shows that in our survey samples, those with above- and below-median time with the card are balanced, consistent with our finding that the treatment timing of localities included in the rollout is not correlated with observables.[17] As a robustness check, we also add controls to specification (4) for the household-level characteristics from Table III, Panel A, and the baseline locality-level characteristics from Figure 3.

It is worth emphasizing that the beneficiaries in the household surveys are a strict subset of the beneficiaries in the administrative data, and that the underlying variation in all specifications stems from time with the card, which was determined exogenously by Oportunidades' rollout of debit cards.

IV. Effect of Debit Cards on Account Use and Savings

In this section, we use administrative data from Bansefi on all transactions and average monthly balances in 348,802 accounts of Oportunidades beneficiaries to estimate the dynamic effect of debit cards on transactions (withdrawals and deposits) and savings in the accounts. To interpret the results, we first note that beneficiaries began using their debit cards to make withdrawals at ATMs almost immediately (rather than continuing to make withdrawals at bank branches). Figure 4 shows that in the four-month period in which they received cards, 83% of beneficiaries withdrew money from an ATM. This percentage increased to around 90% in subsequent periods.

A subset of beneficiaries (45% on average across periods) also used the cards to make purchases at POS terminals. Conditional on making a POS transaction, they averaged 2.2 transactions per period, with the average amount spent per POS transaction equal to 92 pesos (US$7). Throughout the remainder of the paper, "withdrawal" includes bank withdrawals, ATM withdrawals, and POS transactions.

[16] An additional issue with this specification is that, to the extent that treatment had immediate effects, we may be biased against finding an effect since all of our observations in this analysis are treated.

[17] The *Payment Methods Survey* includes fewer measures of household and sociodemographic characteristics since the survey focused on experience with debit cards and ATMs. We find no statistically significant differences in the 10 variables on household and sociodemographic characteristics included in the *Trust Survey* nor the five variables included in the *Payment Methods Survey*.

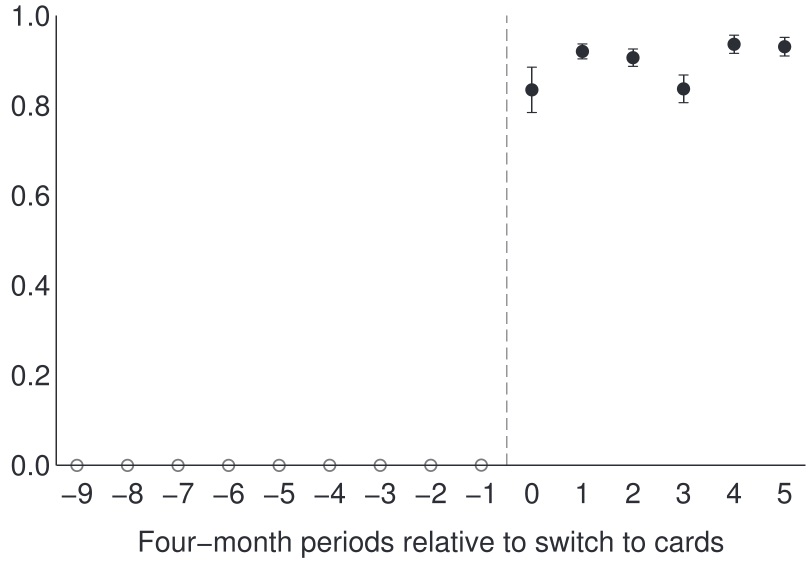

Figure 4. Share of clients using debit cards to withdraw at ATMs. This figure shows the share of clients using their debit card for at least one withdrawal during a four-month period. It shows that beneficiaries immediately adopted the new technology and used their cards to withdraw their transfers, instead of going to the Bansefi bank branch. Note that in periods before receiving the card, the share of clients using debit cards to withdraw at ATMs was necessarily zero. These results are also shown in Table IA.III. N = 2,799,372 account-period observations from 255,781 treated beneficiaries. The dashed vertical line indicates the timing of debit card receipt. Standard errors are clustered at the locality level. Whiskers denote 95% confidence intervals. Black circles indicate results that are significant at the 5% level.

A. Number of Transactions

By reducing indirect transaction costs, debit cards should lead to more transactions, as predicted both by theory (Baumol (1952), Tobin (1956)) and by prior empirical evidence (Attanasio, Guiso, and Jappelli (2002), Schaner (2017)). This is indeed what we find. Figure 5, Panel A, shows the distribution of the number of withdrawals per bimester, before and after receiving the card. Prior to receiving the card, 90% of beneficiaries made a single withdrawal per bimester. The distribution of withdrawals in the control group was nearly identical to that of the treatment group prior to receiving a debit card. In contrast, after receiving the card, 67% of beneficiaries continued to make just one withdrawal, but 25% made two withdrawals, 5% made three withdrawals, and 2% made four or more withdrawals. Meanwhile, the number of withdrawals in the control group did not change over time (Figure IA.3). Recall that the first two withdrawals per bimester were free at any bank's ATM, but subsequent withdrawals were charged a fee; this may explain why few beneficiaries made more than two withdrawals even after receiving the card.

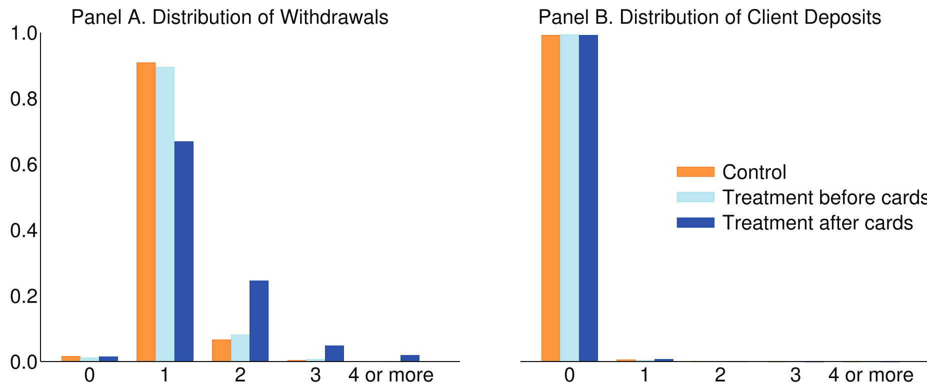

Figure 5. Distribution of withdrawals and client deposits per bimester. This figure shows that after receiving a card, a substantial portion of beneficiaries began making two, three, or four or more withdrawals per bimester rather than one. It depicts the distribution of withdrawals per bimester in Panel A and of client deposits (i.e., excluding Oportunidades deposits) per bimester in Panel B. The three categories represent accounts in the control group, the treatment group before receiving the card, and the treatment group after receiving the card. Within each group, all account-bimester observations are included. $N = 35{,}236{,}129$ transactions from 348,802 beneficiaries over five years. (Color figure can be viewed at wileyonlinelibrary.com)

At the same time, we find no effect on client deposits. Figure 5, Panel B, shows that 99% of accounts had zero client deposits per bimester before and after receiving the card. This result implies that account holders did not add savings from other sources of income to their Bansefi accounts, which is unsurprising for two reasons. First, deposits could not be made at ATMs (since these belonged to banks other than Bansefi). Because beneficiaries still needed to travel to a Bansefi bank branch to make deposits, the transaction cost of depositing money into the account remained unchanged. Second, beneficiaries received about one-fourth of their total income from the Oportunidades program on average, so unless the optimal savings rate in a particular period was higher than 25% of total income, there would have been no reason to deposit more into the savings account from other income sources.

To examine the evolution of the debit card's effect on withdrawals over time, we estimate the generalized difference-in-differences or event study from specification (1) with withdrawals per bimester as the dependent variable. Figure 6, Panel A, plots the ϕ_k coefficients of average withdrawals per bimester for each four-month period, compared to the period just before receiving cards (also shown in Table IV, column (1)). Prior to receiving the card, pretrends are indistinguishable between treatment and control: we cannot reject the null that $\phi_k = 0$ for all $k < 0$. In addition to having parallel trends, both treatment and control accounts averaged just under one withdrawal per period on average. The effect on withdrawals was immediate, as would be expected from the instantaneous change in transaction costs induced by the card. Specifically, prior to receiving the card, beneficiaries in both the treatment and the control

Table IV

Effect of Debit Cards from Administrative Data

This table shows the effect of debit cards on the key outcomes of interest. Columns (1) and (2) show the effect of cards on the number of withdrawals per bimester and the stock of savings for each four-month period relative to the card shock, estimated using specification (1). These results are equivalent to those shown in Figure 6. Columns (3) to (6) show the effect of cards on the number of balance checks by card recipients relative the final period in the sample, estimated using specification (6); they correspond to the four definitions of balance checks used in the paper: all balance checks (column (3)), balance checks after receiving the cash transfer (column (4)), balance checks after the first withdrawal of the bimester (column (5)), and balance checks not in the week prior to a POS transaction (column (6)). These results are equivalent to those shown in Figure 8. Asymptotic standard errors clustered at the locality level are included in parentheses. Randomization inference p-values based on 2,000 draws are included in square brackets. *, **, and *** indicate statistical significance at $p < 0.1$, $p < 0.05$, and $p < 0.01$, respectively (based on asymptotic cluster-robust standard errors).

	Number of Withdrawals	Stock of Savings	Number of Balance Checks			
Period	(1)	(2)	(3)	(4)	(5)	(6)
−9	−0.03	7.08				
	(0.03)	(24.32)				
	[0.28]	[0.79]				
−8	0.00	−38.87*				
	(0.01)	(23.06)				
	[0.93]	[0.13]				
−7	0.06	−19.86				
	(0.07)	(24.31)				
	[0.52]	[0.46]				
−6	0.00	−9.24				
	(0.02)	(14.18)				
	[0.93]	[0.56]				
−5	−0.00	−25.74				
	(0.01)	(16.97)				
	[0.83]	[0.14]				
−4	0.00	−32.58**				
	(0.02)	(12.76)				
	[0.99]	[0.01]				
−3	0.01	−16.37				
	(0.01)	(12.40)				
	[0.33]	[0.21]				
−2	−0.00	4.74				
	(0.01)	(16.67)				
	[0.73]	[0.78]				
−1						
0	0.36***	120.85***	0.87***	0.79***	0.94***	1.06***
	(0.04)	(33.05)	(0.06)	(0.06)	(0.07)	(0.07)
	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
1	0.34***	91.37**	0.81***	0.76***	0.67***	0.90***
	(0.03)	(44.29)	(0.05)	(0.04)	(0.03)	(0.03)
	[0.00]	[0.03]	[0.00]	[0.00]	[0.00]	[0.00]
2	0.30***	125.09***	0.45***	0.49***	0.51***	0.62***
	(0.02)	(30.99)	(0.05)	(0.04)	(0.04)	(0.05)
	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]

(*Continued*)

Table IV—*Continued*

	Number of Withdrawals	Stock of Savings	Number of Balance Checks			
Period	(1)	(2)	(3)	(4)	(5)	(6)
3	0.33***	447.48***	0.35***	0.08***	0.09***	0.16***
	(0.02)	(41.98)	(0.03)	(0.01)	(0.01)	(0.02)
	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
4	0.24***	661.54***	0.21***	0.04***	0.02**	0.04***
	(0.04)	(56.68)	(0.02)	(0.01)	(0.01)	(0.01)
	[0.00]	[0.00]	[0.00]	[0.02]	[0.03]	[0.01]
5	0.26***	767.87***				
	(0.03)	(56.72)				
	[0.00]	[0.00]				
N observations	4,740,331	4,668,575	873,848	873,848	873,848	873,848
N accounts	348,802	348,802	233,080	233,080	233,080	233,080

groups averaged about one withdrawal per bimester, but immediately after receiving the card, treated beneficiaries began making an additional 0.36 withdrawals per bimester on average. Table IA.II shows that the results are not sensitive to winsorization, to including baseline account characteristics interacted with time dummies, or to using the inverse hyperbolic sine transformation for the outcome.

B. The Stock of Savings (Account Balances)

Next, we explore whether debit cards caused an increase in savings from period to period. The results on withdrawals tell us nothing about period-to-period savings. On the one hand, beneficiaries could have continued withdrawing once per period but reduced the total amount withdrawn during the period, thereby increasing their savings. On the other hand, they could have increased the number of withdrawals and left some money in the account after the first withdrawal in the pay period, but withdrawn the remaining money later in the same period, thereby leaving the account balance close to zero by the end of that period. The latter possibility means that we have to construct our measure of savings carefully, as this behavior would lead to a mechanically higher average balance *within* each period that does not correspond to accumulating savings in the account over time, that is, *across* periods.

Instead, we are interested in measuring savings across periods. Ideally, we would measure the stock of savings as the end-of-period balance, calculated as the beginning-of-period balance plus all deposits minus all withdrawals. Unfortunately, while we have transactions data beginning in January 2007, we do not have the initial balances for the first period of our data (January 2007), and cannot assume that these equal zero because we know from the average balance data that a nontrivial share of beneficiaries saved in their accounts prior to 2007. Thus, to construct a reliable estimate of end-of-period balance,

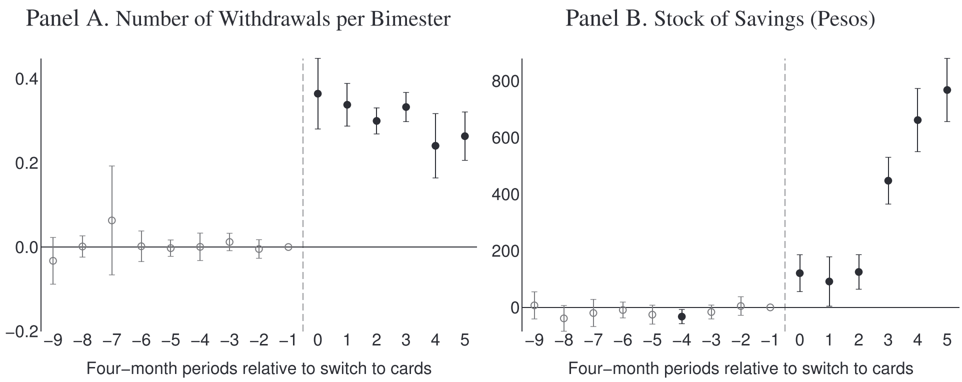

Figure 6. Effect of debit cards on withdrawals and savings. Panel A shows the effect of debit cards on the number of withdrawals per bimester. The figure plots the ϕ_k coefficients from specification (1), where the dependent variable is number of withdrawals. $N = 4{,}740{,}331$ account-period observations from 348,802 beneficiaries. Panel B shows the effect of debit cards on the stock of savings. The figure plots the ϕ_k coefficients from specification (1), where the dependent variable is the end-of-period account balance. $N = 4{,}664{,}772$ account-period observations from 348,802 beneficiaries. The dashed vertical lines indicate the timing of debit card receipt. Standard errors are clustered at the locality level. Whiskers denote 95% confidence intervals. Black circles indicate results that are signficant at the 5% level, while hollow circles indicate results that are statistically insignificant from zero.

we combine data on the average balance for each period with transaction-level data on the timing and amount of each transaction (see Internet Appendix Section IV for details).

We estimate specification (1) using account i's end-of-period balance in period t as the dependent variable. Following other papers measuring savings (e.g., Kast, Meier, and Pomeranz (2018)), we winsorize savings balances at the 95^{th} percentile to avoid results driven by outliers. The ϕ_k terms thus measure the causal effect of debit cards on the stock of savings k periods after receiving a card. Figure 6, Panel B, plots the ϕ_k coefficients and their 95% confidence intervals (also shown in Table IV, column (2)). We note the empirical support for parallel trends, shown by the zero coefficients for pre-event periods, $k < 0$.[18] In the first few periods after receiving a card, we observe a small savings effect of 100 to 200 pesos (about US\$8 to US\$15). The initial effect is small because only a minority of beneficiaries began saving shortly after receiving a card; we explore this further below. Savings increased substantially after about one year with the card—three four-month periods after card receipt, the savings effect is 447 pesos, while it is 768 pesos after two years with the card. These effect sizes are equal to 1.2% and 2.0% of annual income, respectively, and are larger than the effect sizes found in other studies of savings interventions (Figure 1). Table IA.II shows that the results are not sensitive to winsorization, to

[18] In eight of the nine pretreatment periods, there is no statistically significant difference between the savings balance of the treatment and control groups.

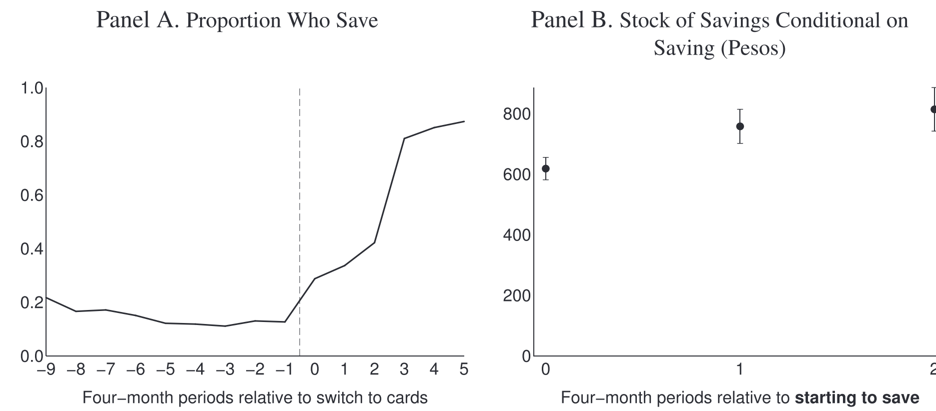

Figure 7. Decomposition of savings effect. This figure decomposes the effect of debit cards on the stock of savings into the extensive margin effect on the proportion who saved over time, and the intensive margin effect on the stock of savings conditional on saving. Panel A shows the proportion of treated beneficiaries who saved in each period relative to when they received a debit card. These results are also shown in Table IA.III. N = 2,968,628 account-period observations for 255,784 treated beneficiaries. The dashed vertical line indicates the timing of debit card receipt. Panel B plots ϕ_k from specification (1) with the event time dummies redefined relative to when an individual started saving in the account, and we impose a zero pretreatment trend by setting $a = 0$. These results are also shown in Table IA.V. N = 4,668,575 account-period observations from 348,802 beneficiaries. Standard errors are clustered at the locality level. Whiskers denote 95% confidence intervals. Black circles indicate results that are significant at the 5% level.

including baseline locality characteristics interacted with time dummies, or to using the inverse hyperbolic sine transformation for the outcome.

The effect of debit cards on the average stock of savings shown in Figure 6, Panel B, combines two effects: the impact of debit cards on the probability of saving and the savings amount conditional on saving. Figure 7 decomposes these two components. Panel A shows the proportion of treated beneficiaries who had at least a small positive balance at the end of each period: while only 13% of beneficiaries saved in their account in the period before receiving cards, an additional 16% started saving immediately after receiving a card.[19] For these beneficiaries, it is likely that the reduction in the transaction costs of accessing savings provided by the cards was a sufficient condition to save in a formal bank account. The proportion of beneficiaries who saved in their Bansefi accounts increased over time: after nearly one year with the card, 42% of beneficiaries saved in the account, and after two years nearly all beneficiaries (87%) saved in their Bansefi account.

To estimate the second component, the amount of savings conditional on having started to save, we define a new event as the period in which a beneficiary

[19] We set the threshold for a "small positive balance" at 150 pesos, given that balances below 50 or 100 pesos could have been due to ATMs not disbursing exact change rather than voluntary savings.

began saving (rather than when the beneficiary received a card). This event is not causal—it occurred at different points in time for different beneficiaries, due to both the quasi-exogenous timing of receiving cards and the endogenous timing of when they chose to start saving once they received the card. Our goal with this estimation is merely descriptive—to estimate the amount of savings each period after having started to save. We estimate specification (1) using this new event and plot the results in Figure 7, Panel B.[20] In the first period when beneficiaries saved in the account, they deposited 618 pesos on average, or 4.9% of their total income that period. They deposited significantly less in the following periods, consistent with models of precautionary saving in which an individual's savings rate is decreasing in her stock of savings as it approaches her buffer stock target (Carroll (1997)).

V. Increase in Overall Savings versus Substitution

The increase in formal savings in beneficiaries' Bansefi accounts might represent a shift from other forms of saving, such as saving under the mattress or in informal saving clubs, with no change in overall savings. This section investigates whether the observed increase in Bansefi account savings was an increase in overall savings or crowded out other savings. We take advantage of Oportunidades' *Household Panel Survey*, conducted in four waves during the years 2002, 2003, 2004, and November 2009 to February 2010.

We use a simple difference-in-differences identification strategy where we examine changes in beneficiaries' consumption, income, and stock of assets, again exploiting the differential timing of debit card receipt. We compare those with cards at the time of the fourth survey wave to those who had not yet received cards. Section III formally tested for parallel pretreatment trends for each dependent variable and failed to reject the null hypothesis of parallel trends. We estimate specification (2) in Table V, columns (1) to (3), and specification (2) augmented with the additional interaction of time fixed effects and baseline household characteristics in column (4), separately for three dependent variables: consumption, income, and an asset index.[21] The additional interaction of time fixed effects and baseline household characteristics follows de Janvry et al. (2015) and absorbs variation in how the dependent variable evolved over time for different types of households.

[20] Because the majority did not begin saving until they had the card for a year, we only graph the savings stock for three postsaving periods (as further-period estimates would be based solely on the small sample of earlier savers).

[21] Standard errors shown in parentheses are cluster-robust asymptotic standard errors, clustered at the locality level. There are 46 localities in our estimation sample from the survey. We also show wild cluster bootstrap percentile-t 95% confidence intervals in square brackets, as well as clustered randomization inference p-values in square brackets.

Table V

Effect of Debit Cards from Household Panel Survey

This table shows the effect of debit cards on consumption, income, and assets using the *Household Panel Survey* combined with administrative data from Oportunidades on the debit card rollout, estimated using specification (2). Each row label is the dependent variable from a separate regression; each column is a different specification. Means for each dependent variable can be found in Table III, Panel B. Dependent variables are measured in pesos per month, with the exception of the asset index. The asset index is the first principal component of assets that are included in both the pretreatment (2002, 2003, 2004) and posttreatment (2009 to 2010) versions of the survey: car, truck, motorcycle, television, video or DVD player, radio or stereo, washer, gas stove, and refrigerator. The number of households in column (4) is lower because households have missing values for one of the household characteristics included, or are not included in enough pretreatment waves to construct household-level pretrends of the outcome variables, which are interacted with time fixed effects in that specification. Asymptotic cluster-robust standard errors (clustered at the locality level, using pretreatment locality) are included in parentheses. Wild cluster bootstrap percentile-t 95% confidence intervals based on 1,000 draws are included in square brackets. Randomization inference p-values based on 2,000 draws are included in square brackets. *, **, and *** indicate statistical significance at $p < 0.1$, $p < 0.05$, and $p < 0.01$, respectively (based on asymptotic cluster-robust standard errors).

	(1)	(2)	(3)	(4)
Consumption	−175.36**	−150.51**	−136.52**	−155.11**
	(81.31)	(70.43)	(61.75)	(62.07)
	[−353.11, −1.52]	[−306.24, −2.30]	[−276.37, −4.75]	[−288.02, −33.10]
	[0.04]	[0.04]	[0.04]	[0.02]
Income	98.16	106.01	75.50	38.11
	(170.03)	(150.31)	(127.77)	(106.12)
	[−290.77, 486.11]	[−230.64, 468.97]	[−219.75, 376.72]	[−175.00, 251.64]
	[0.63]	[0.56]	[0.61]	[0.74]
Asset index	0.06	0.06	0.07	0.03
	(0.08)	(0.08)	(0.07)	(0.08)
	[−0.12, 0.24]	[−0.12, 0.24]	[−0.08, 0.23]	[−0.20, 0.24]
	[0.54]	[0.54]	[0.41]	[0.82]
p-Value consumption vs. income	[0.047]	[0.041]	[0.056]	[0.057]
Number of observations	9,246	9,246	9,246	7,754
Number of households	2,868	2,868	2,868	2,200

A. Total Consumption and Income

Table V, column (4), shows that consumption decreased by about 155 pesos per month among treated households relative to control (statistically significant at the 5% level). We do not find any effect on income. We also test the difference in coefficients of consumption and income using a stacked regression (which is equivalent to a seemingly unrelated regression when the same regressors are used in each equation, as is the case here); although both consumption and income are noisily measured, the difference in the coefficients is significant at the 5% or 10% level in all specifications (the p-value of the F-test of equality of the coefficients on consumption and income is 0.057 in column (4)). Table V, columns (1) to (3), show that our results are robust to the extent of winsorizing and to removing the controls for flexible time trends as a function of household characteristics.

The point estimates of the effect of debit cards on consumption and the lack of an effect on income suggest that the increase in formal savings shown in Section IV represents an increase in total savings (since total savings is income minus consumption). To compare estimates from the survey and administrative data, we first note that the survey estimate on reduced consumption—which equals 4.9% of income—is measured as a flow at a specific point in time relative to receiving a card. Specifically, this survey estimate corresponds to the effect of a card on the flow of consumption after approximately one year.[22] The timing of this effect matters for our comparison given our finding from the administrative data that the savings stock did not evolve linearly over time.

We therefore compare our estimate of the reduction in consumption from the survey (4.9% of income) to the change in the flow of savings after one year with the card, for a comparable set of beneficiaries in the administrative data. We achieve this by restricting the administrative data for this comparison to the same set of treatment localities included in our survey estimates. Recall that we cannot restrict the analysis to the exact same accounts because we cannot merge the two data sources at the account/beneficiary level. We then compute the change in the average $\Delta Savings_{it}$ for this subset of the administrative data, where we restrict t to the period one year after receiving the card. For ease of interpretation, we divide this change in the flow of savings after one year with the card, measured in pesos, by average income (taken from the survey).[23] This gives us an estimate of the effect of debit cards on the flow of formal savings of 4.6% of income, which is within 0.2% of income—or less than 50 cents per

[22] This is because the treated localities included in the posttreatment survey wave were deliberately selected to primarily be localities that received cards at the beginning of the rollout; see Figure IA.1. We exclude beneficiaries in the survey who lived in localities that received cards shortly before the survey wave.

[23] As usual, Δ is computed relative to the preceding four-month period. The change in $\Delta Savings_{it}$ is relative to the period before receiving cards (when, in any event, beneficiaries were not accumulating savings in their accounts). Since the change in the flow of savings in pesos is calculated for a four-month period, we convert monthly income from the survey to four-month income before dividing.

month—of our survey estimate of reduced consumption. Furthermore, each of the two estimates is within the 95% confidence interval of the other.

As in most household surveys, however, our estimates of the change in consumption are noisy: while we can reject the hypothesis that the increase in formal savings was purely substitution from other forms of saving, we cannot rule out the possibility that some but not all of the increase in formal savings was substitution. To maximize power, we focus on the specification from Table V, column (4), which absorbs additional variation in how the dependent variable evolved over time for different types of households by including time fixed effects interacted with baseline household characteristics. The lower bound of the 95% confidence interval estimated using a percentile-t wild cluster bootstrap is a reduction in consumption equal to 33 pesos per month or 1.0% of income; the lower bound of the 90% confidence interval is 52 pesos per month or 1.6% of income.

B. Exploring the Fall in Consumption

The finding that consumption decreased by the same amount as formal savings increased, while income did not appear to change, suggests that the increase in formal savings represents new savings rather than purely substitution from other forms of saving. Why would receiving debit cards increase overall savings, financed by a reduction in consumption? One possibility is that cash is "hot" in hand (or when being saved at home) and that it is easier for other household members to access the money when saved at home rather than in a bank account (Ashraf (2009)).

Under this hypothesis, receiving a card should cause consumption to fall relatively more in categories where temptation is the greatest. We test and find suggestive evidence in support of this prediction in Table VI, which shows the effect of debit cards on the proportion of income spent on various consumption categories (estimated using specification (2) with the proportion of income spent on a category as the dependent variable). To assess the relative change in consumption in each category, column (3) shows the point estimates divided by the control group's proportion of income spent on that category. We find a more negative point estimate (−14%) for the change in consumption of temptation goods than for other consumption.[24] However, households also reduced consumption in other categories: consumption of other food and drink and of other nondurable goods (clothing, personal care, household cleaning items, and fuel) changed by −10% each. Although we cannot reject the possibility that the coefficient on temptation goods equals those of other food and drinks or other nondurable goods, we can reject the hypothesis that it equals the positive but not statistically significant coefficient on education and health spending.

We also use the survey to test whether the increase in formal savings observed in the administrative bank account data crowded out a particular form

[24] Temptation goods are defined based on Banerjee and Mullainathan (2010), and include alcohol, tobacco, sugar, soda, sweets, junk food, and fats.

Table VI

Effect of Debit Cards on Proportion of Income Spent by Category

This table shows the effect of debit cards on consumption by category, estimated using specification (2), where the outcome variable is the proportion of income spent on each category. Column (1) shows the control group's mean proportion of income spent on each category at baseline to show the relative importance of the categories in total consumption. Column (2) shows the coefficients from specification (2). Column (3) divides those coefficients by the control group's mean proportion of income spent on that category of consumption to show the relative change in the proportion of income spent on each category. We include household baseline characteristic × time fixed effects to increase power, as in our preferred specification in Table V, column (4). The number of households and observations is lower than in Table V, column (4), due to missing values in particular consumption categories. Asymptotic standard errors clustered at the locality level are included in parentheses. Randomization inference p-values based on 2,000 draws are included in square brackets. *, **, and *** indicate statistical significance at $p < 0.1$, $p < 0.05$, and $p < 0.01$, respectively (based on asymptotic cluster-robust standard errors).

	Control Baseline Mean (1)	Effect on Proportion of Income (2)	Relative Change (3)	N (4)	Number of Households (5)
Temptation goods	0.082	−0.017***	−0.138***	7,077	2,066
	(0.003)	(0.006)	(0.047)		
		[0.011]	[0.011]		
Other food and drinks	0.557	−0.086**	−0.104**	7,077	2,066
	(0.017)	(0.034)	(0.042)		
		[0.065]	[0.065]		
Other nondurable goods	0.151	−0.023**	−0.096**	7,077	2,066
	(0.008)	(0.009)	(0.037)		
		[0.031]	[0.031]		
Education and health	0.070	0.008	0.095	7,077	2,066
	(0.003)	(0.006)	(0.069)		
		[0.285]	[0.285]		
Other services	0.001	−0.000	−0.073	7,077	2,066
	(0.000)	(0.000)	(0.139)		
		[0.656]	[0.656]		

of informal saving, namely, investment in durable assets. We test whether beneficiary households disinvested or invested less in assets by constructing an asset index. We find that the difference-in-differences coefficients on this measure are small, positive, and statistically insignificant (Table V). Together with the results on education and health spending, this suggests that beneficiaries were not increasing savings by substituting from investments in human or physical capital, but rather by decreasing their consumption of nondurable goods. We note, however, that our time horizon might be too short to see potentially positive effects of saving on future investments.

Despite the prevalence of ATMs and POS terminals, which make it easy to spend money saved in a bank account, the debit card likely increased savings for two reasons. First, saving in the bank account prior to having a debit card involved high transaction costs and beneficiaries had low trust in the bank, both of which prevented formal savings. Second, intrahousehold bargaining issues could make saving informally at home difficult when household members have different preferences (Anderson and Baland (2002), Schaner (2015)). In other words, it may have been difficult for the women receiving transfers to save at home due to a lack of control over their partners' access to the savings. To test this hypothesis, we construct an index of bargaining power at baseline for households who had at least one male adult present in addition to the female cash transfer beneficiary.[25] Results in Table IA.VI suggest that the effect of debit cards on consumption (and hence on savings) was concentrated among households in which the woman had below-median baseline bargaining power. We note that while the coefficients on the interaction term are quantitatively large, they are only marginally significant in some specifications and insignificant in others, and hence, these results are speculative.

VI. Mechanisms

The card decreased indirect transaction costs to both access savings and monitor account balances. In this section, we provide evidence that both mechanisms were at work in increasing active account use and savings. We also explore other potential mechanisms such as learning the ATM technology.

A. Transaction Costs to Access Account

Consistent with economic theory on the effect of an immediate decrease in transaction costs (Baumol (1952), Tobin (1956)), we observe an immediate increase in the number of withdrawals per period (Figure 6, Panel A). The percentage of clients who used their debit card to make at least one withdrawal at an ATM instead of going to the bank branch also increased immediately after receiving the card—to 83% of beneficiaries—after which it was fairly stable (Figure 4). We also observe that 16% of beneficiaries began saving immediately after receiving the card, likely due to the change in transaction costs (Figure 7, Panel A).

While we observe an immediate increase and then flat time profile of the share of beneficiaries withdrawing their benefits at ATMs (Figure 4) and the number of withdrawals per period (Figure 6, Panel A), the share of beneficiaries who started saving in their Bansefi accounts increased only gradually over time after receiving cards (Figure 7, Panel A). This gradual increase could be explained in part by the transaction costs of accessing the account or by the gradual increase in retailer adoption of POS terminals documented in Higgins (2019). Because these mechanisms likely explain only part of the gradual increase in the proportion of beneficiaries who saved—for example, for the POS terminal expansion, after two years with the card 87% of beneficiaries saved in the account but only 44% used the card to make POS transactions—we investigate other mechanisms beyond transaction costs to access money in the

[25] The bargaining power index is based on five questions about whether important household decisions are made by the woman, the man, or jointly. Details are in the notes to Table IA.VI.

account. In particular, in the remainder of this subsection, we test whether other types of transaction costs changed over time; in the next subsection, we investigate beneficiaries' use of the debit card to monitor their account and build trust in the bank over time.

First, we test and rule out the possibility that banks disproportionately expanded complementary infrastructure (e.g., the number of ATMs) in treated localities, which would have further decreased the transaction cost of accessing funds in a way that is geographically correlated with the debit card expansion. We use quarterly data on the number of ATMs and bank branches by municipality from CNBV for the period between the last quarter of 2008—the first quarter with available data—through the last quarter of 2011. We estimate a difference-in-differences specification with six leads and lags,

$$y_{mt} = \lambda_m + \delta_t + \sum_{k=-6}^{6} \beta_k D_{m,t+k} + \varepsilon_{mt}, \tag{5}$$

where y_{mt} is the number of total ATMs, total bank branches, Bansefi ATMs, or Bansefi branches in municipality m in quarter t, and D_{mt} is an indicator equal to 1 if at least one locality in municipality m had Oportunidades debit cards in quarter t. We conduct an F-test of whether lags of debit card receipt predict banking infrastructure (i.e., whether there was a supply-side response by banks to the rollout of debit cards: $\beta_{-6} = \cdots = \beta_{-1} = 0$), and an F-test of whether leads of debit card receipt predict banking infrastructure (i.e., whether debit cards were first rolled out in municipalities with a recent expansion of banking infrastructure: $\beta_1 = \cdots = \beta_6 = 0$). We find evidence of neither relationship (Table IA.VII).

Second, we test whether the increase in the proportion of savers over time with the card could be explained by a concurrent increase in the number of ATMs across all localities. Only beneficiaries in treatment localities could access money at ATMs and hence take advantage of an expansion of ATMs. If the gradual increase in the proportion saving over time were due to a gradual decrease in transaction costs that was uncorrelated with the geographical expansion of debit cards, we would also expect savings to increase among Bansefi debit card holders who were not Oportunidades beneficiaries. We look at mean savings among non-Oportunidades debit card account holders who opened their accounts in 2007 and hence had the account for about two years when our study period began. Figure IA.4 shows that savings among non-Oportunidades debit card holders did not increase over the study period, but instead stayed relatively flat. This result suggests that the increase over time in the proportion that saved cannot be explained by a gradual decrease in transaction costs over time.

Third, beneficiaries' perceptions of transaction costs might have changed over time, even if transaction costs associated with the card remained constant. For example, perhaps beneficiaries were checking balances to learn about direct transaction costs (i.e., fees), in which case they would check balances less frequently after learning these transaction costs. We directly test and rule out this possibility using the *Payment Methods Survey*, which asked beneficiaries how much the bank charges them for (i) a balance check and (ii) a withdrawal after the initial free withdrawals. We find that beneficiaries tended to get the level of these fees right and, more importantly, there is no difference across beneficiaries who had the card for above- versus below-median time (Table IA.VIII).

B. Monitoring Costs and Trust

A lack of trust in banks is frequently cited by the poor as a primary reason for not saving (Dupas et al. (2016), FDIC (2016)). The time delay between receiving the debit card and starting to save (for most beneficiaries) is consistent with the hypothesis that the debit card reduced the indirect cost of checking account balances, leading to an increase in balance checks to monitor that the bank was not regularly reducing beneficiaries' account balances.

Under this hypothesis, each additional balance check would provide additional information about the bank's trustworthiness. With simple Bayesian learning, balance checks would have a decreasing marginal benefit as a beneficiary updated her beliefs about the bank's trustworthiness, which would lead to a decrease in the number of balance checks over time. Hence, over time with the card, we expect the number of balance checks to fall and trust to rise. Below, we use administrative data to test whether balance checks fell over time and survey data to test whether self-reported trust in the bank increased over time with the card.

B.1. Balance Checks Fell over Time with the Debit Card

We first use the Bansefi transactions data to test whether balance checks fell over time with the card. We only observe balance checks once beneficiaries had debit cards, which restricts our analysis to the treatment group and to periods after the card was received. Pooling data across periods after beneficiaries received cards, we find that on average beneficiaries checked their balances 1.7 times per four-month period. To test the hypothesis of a decreasing time trend in balance checking, we regress the number of balance checks on account fixed effects and event time dummies, omitting the dummy for the last period:

$$Balance\ Checks_{it} = \lambda_i + \sum_{k=0}^{\overline{k}_i - 1} \pi_k D_{it}^k + \varepsilon_{it}. \tag{6}$$

The π_k coefficients estimate the number of balance checks k periods after receiving the card relative to the last period in the sample (July to October 2011), which depending on the beneficiary corresponds to one to two years after receipt of the card.[26]

[26] $\overline{k}_i$ denotes the last period with the card for account i in our data, which varies depending on when i received a card. We do not include time fixed effects since we can only include treated

Figure 8, Panel A, plots the π_k coefficients using any balance check to construct the dependent variable, and shows that the number of balance checks in the periods following receipt of the debit card was higher than in later periods (also shown in Table IV, column (3)). For example, in the period after receiving the card, beneficiaries made an average of 0.9 more balance checks compared to two years after receiving the card. After having the card for about one year, this fell to 0.4 more balance checks. For learning to occur, beneficiaries would need a positive balance in their account at the time of checking. We find that in the four months after getting the card, 89% of accounts had a positive (small) balance at the time of a balance check after receipt of the transfer: the 25$^{\text{th}}$ percentile of balances at the time of a balance check is 20 pesos, the median is 55 pesos, and the 75$^{\text{th}}$ percentile is 110 pesos.[27]

Although beneficiaries were given calendars with exact transfer dates and should have known the dates on which transfers were deposited (see Figure IA.10), we additionally use two more restrictive definitions of a balance check, to ensure that a balance check constituted bank monitoring rather than checking whether the Oportunidades deposit had arrived. The first alternative definition excludes all balance checks that occurred prior to the transfer being deposited that bimester, since these checks might have been made to verify whether the transfer had arrived. This definition also excludes balance checks that occurred on the same day as a withdrawal, because if a beneficiary was checking whether the transfer had arrived and she found that it had, she likely would have withdrawn it that day. An even more conservative definition only includes balance checks that occurred after the bimester's transfer had arrived and the client had already made a subsequent withdrawal. Because the next transfer would not arrive until the following bimester, and the beneficiary already made a withdrawal after the transfer arrived in the current bimester, the beneficiary would have known that the current bimester's transfer had arrived. Hence, these checks could not represent an attempt to verify whether the transfer had arrived. Figure 8, Panels B and C, plots the results using these two alternative definitions and show a very similar decrease in balance checks over time (also shown in Table IV, columns (4) and (5)).

A separate possibility is that beneficiaries used balance checks to ensure that they had money in their account before making a transaction at a POS terminal. Bansefi did not charge overdraft fees—if a beneficiary attempted to make a purchase at a POS terminal but did not have enough money in the account, the transaction was denied with an "insufficient funds" message sent to the POS terminal. However, while the beneficiary would not face a

beneficiaries after treatment in the regression, and the within-account trend in balance checks over time (among this group) is precisely the variation we are exploiting. Standard errors are clustered at the locality level.

[27] For these statistics, because we do not have initial January 2007 balance (and hence do not know the precise balance at any point in time), we take the conservative approach of defining a balance as positive if the cumulative transfer amount minus the cumulative withdrawal amount in the bimester was positive at the time of the balance check. This is a sufficient but not necessary condition for the balance to be positive.

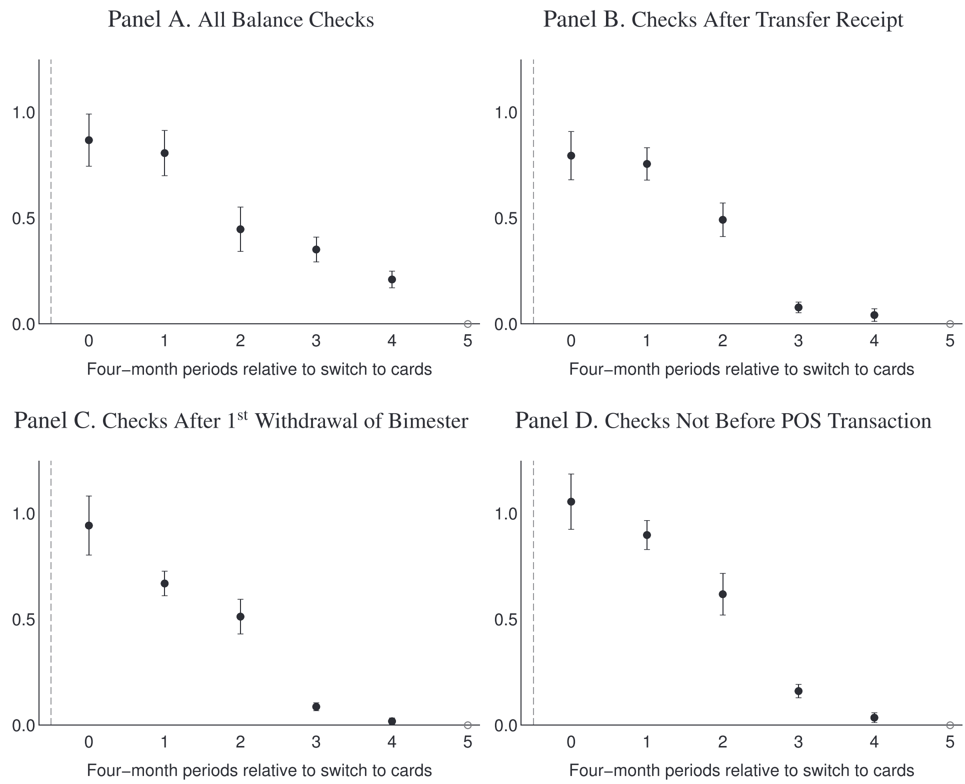

Figure 8. Number of balance checks over time. This figure shows the number of balance checks over time after receiving debit cards. It shows the number of balance checks relative to the last period in the data for each observation by plotting the π_k coefficients from specification (6). Periods before receiving the card are not included since it was only possible to check balances at Bansefi branches without a debit card, and these balance checks are not recorded in our data. Panel A includes all balance checks, while Panels B to D correspond to a narrower definition of balance checks, where the narrower definitions attempt to rule out balance checks for purposes other than monitoring the bank. Panel B shows balance checks after the transfer was received and on a different day than a withdrawal, Panel C shows balance checks after the first withdrawal occurred in the bimester and on a different day than a withdrawal, and Panel D shows balance checks not within the seven days before or the day of a POS transaction (see Section VI.B.1). $N =$ 873,848 account-period observations from 233,080 unique treated beneficiaries. Accounts in which cards were received in the last period of our data must be excluded to omit a D_{it}^{k} dummy; we also exclude beneficiaries who received the card in the second-to-last period in our data since we only observe one additional postcard period for those beneficiaries. Standard errors are clustered at the locality level. Whiskers denote 95% confidence intervals. Black circles indicate results that are significant at the 5% level.

monetary penalty for attempting to make a debit card purchase with insufficient funds, there might be a social penalty: the beneficiary may prefer to avoid having their transaction denied when attempting to make a purchase. Indeed, some balance checks do appear to have been made to ensure that money was in the account prior to making a transaction, as the number of balance checks increased in the seven days preceding or day of a POS transaction (Figure IA.5). However, excluding checks made on the same day as an ATM withdrawal, only 20% of balance checks were in the week preceding a POS transaction. Furthermore, Figure 8, Panel A, shows that the magnitude of the decrease in balance checks over time is similar when using an alternative definition of balance checks that excludes checks made in the week preceding a POS transaction (also shown in Table IV, column (6)).[28]

We validate the above results using survey data from the *Payment Methods Survey*. Specifically, we estimate specification (4) using the self-reported number of balance checks (without withdrawing cash) over the past four-month period as the dependent variable. Table IA.VIII shows that those who had the card for more than the median time (12 months) reported making 31% fewer trips to the ATM to check their balances without withdrawing money than those who had the card for less time. The self-reported survey responses thus confirm the findings from the administrative data, and also show that balance checking behavior was salient for beneficiaries.

B.2. Trust Increased over Time with the Debit Card

We next test the hypothesis that longer tenure with the debit card induced higher trust in the bank. As described in Section II.B, the *Trust Survey* first asked the beneficiary if she saves in her Bansefi bank account, and if she answered no, it asked why not. If she did not save in the account and indicated that she did not trust the bank, we code lack of trust as one; otherwise (including if the beneficiary saved in the account), we code lack of trust as zero.

We estimate specification (4) using lack of trust as the dependent variable, again exploiting the exogenous variation in the length of time beneficiaries had the card. As explained in Section III, to interpret γ in specification (4) as a causal effect, we need to assume that time with the card is orthogonal to our potential outcomes of interest. The balance tests conducted for the *Trust Survey* sample support this assumption (Table III, Panel A), as does the finding that conditional on being included in the debit card rollout, the timing of when cards were distributed to the locality is uncorrelated with observables (Table II and Figure 3). Table VII shows that trust increased over time: beneficiaries with above-median time with the card were 33% less likely to report not saving due to low trust.[29] For comparison, Table VII also shows results for two alternative forms of learning discussed in Section VI.C: learning to use the technology and learning that the program will not drop beneficiaries

[28] Figure IA.6 illustrates the four definitions of balance checks that we use.

[29] Note that because of the timing of the *Trust Survey*, those with the card for less than the median time still had the card for at least nine months, meaning that some of them would have likely developed trust in the bank prior to being surveyed. Those with above-median time with the card had it for five months longer on average. If anything, this may bias our results downward relative to what we would find if it were possible to compare those who had sufficient tenure with the card to those who had not yet received the card.

Table VII

Self-Reported Reasons for Not Saving in Bansefi Account

This table compares reasons for not saving in the Bansefi bank account among beneficiaries with above- versus below-median time with a debit card, estimated using specification (4). It compares the proportion of respondents in each group who provide the corresponding reason for not saving in response to the questions "Do you leave part of the monetary support from Oportunidades in your bank account?," and if not, "Why don't you leave part of the monetary support from Oportunidades in your Bansefi savings account?" Beneficiaries who reported saving are coded as zero for each reason for not saving and still included in the regressions. Column (1) shows the mean for beneficiaries with below-median time with the card (α) and columns (2) to (4) show the difference in means for those with above-median time with a card (relative to those with below-median time with a card; γ). Column (2) does not include any additional controls. Column (3) controls for the household-level controls that would not be affected by treatment from Table III (number of household members and age, gender, marital status, and education level of the beneficiary). Column (4) controls for both household-level controls and locality- or municipality-level controls for the variables from Figure 3 (log wage, log food prices, log POS terminals, log bank branches, log ATMs, log debit and credit cards, average stock of savings, average log stock of savings, and average number of withdrawals). Asymptotic standard errors clustered at the locality level are included in parentheses. Randomization inference p-values based on 2,000 draws are included in square brackets. *, **, and *** indicate statistical significance at $p < 0.1$, $p < 0.05$, and $p < 0.01$, respectively (based on asymptotic cluster-robust standard errors).

	Mean Card < Median Time	Difference Card ≥ Median Time			N
	(1)	(2)	(3)	(4)	(5)
Lack of trust	0.238	−0.078**	−0.073**	−0.073**	1,694
	(0.033)	(0.033)	(0.030)	(0.031)	
		[0.064]	[0.057]	[0.058]	
Lack of knowledge	0.014	0.008	0.011*	0.011	1,694
	(0.006)	(0.007)	(0.006)	(0.007)	
		[0.340]	[0.178]	[0.245]	
Fear of program ineligibility	0.030	−0.013	−0.013	−0.013	1,694
	(0.009)	(0.009)	(0.011)	(0.011)	
		[0.204]	[0.264]	[0.268]	
Household-level controls		No	Yes	Yes	
Locality/municipality-level controls		No	No	Yes	

who accumulate savings. Few beneficiaries reported these as reasons for not saving, and the proportion did not change over time with the card.

Finally, although we cannot causally relate decreased balance checks with higher account balances, in Figure IA.7, we show that within account, the number of balance checks and savings balances is strongly negatively correlated: in a period when beneficiaries made one balance check, they saved 300 pesos less than in a period when they made no balance checks; when making three or more balance checks, they saved 400 to 500 pesos less.[30] This is consistent

[30] This trend is robust to the definition of a balance check; the numbers we cite here use the two most restrictive definitions of balance checks from Panels C and D. Interestingly, using data from an online financial platform in Iceland, Olafsson and Pagel (2017) find the opposite correlation:

with the mechanism that monitoring balances led to increased trust, which, in turn, led to increased savings over time.

C. Learning

The time delay for many beneficiaries between getting the card and saving suggests some type of learning. Monitoring the bank and building trust is one type of learning. In this section, we explore whether other types of learning occurred. We find no evidence of these other types of learning.

C.1. Learning the Technology

Learning how to use the technology would have to occur gradually over time to explain our results. However, in addition to the survey evidence against this form of learning that we present below, learning the technology is inconsistent with the result from the administrative data that the number of withdrawals and use of ATMs increased *immediately* after receiving the card and remained fairly stable thereafter.

Beneficiaries could have learned how to use their debit cards over time. The *Payment Methods Survey* asked each respondent whether (i) it is hard to use the ATM, (ii) she gets help using the ATM, and (iii) she knows her PIN. We use these three questions as dependent variables in specification (4). Table IA.VIII shows that there is no statistically significant difference between groups who had the card for above- versus below-median time. Beneficiaries could have instead learned how to save in the account (rather than how to use the card). This is unlikely as these beneficiaries already had the account for years prior to receiving a debit card. Consistent with this argument, fewer than 2% of respondents to the *Trust Survey* cited not saving due to lack of knowledge.[31] Moreover, there is no difference between those who had the card for above- versus below-median time (Table VII).

C.2. Learning the Program Rules

Beneficiaries may have initially thought that saving in the account would make them ineligible for the program, but learned over time that this was not the case. In the *Trust Survey*, there are some responses along these lines such as "because if I save in the account, they can drop me from Oportunidades." We thus estimate specification (4) with the dependent variable set to one if respondents did not save for this reason. We find that fewer than 4% of beneficiaries did not save due to fear of being dropped from the program, with no difference between those who had the card for above- versus below-median time (Table VII).

when balances were high, people logged into their accounts more often. The paper then presents a model of anticipatory utility that can explain these findings.

[31] Examples of responses coded as lack of knowledge are "I don't know how to use the card so I withdraw everything at once" and "I don't know how [to save in the account]."

C.3. Time with the Bank Account

Experience with the savings account rather than time with the debit card itself cannot explain the delayed savings effect. First, savings accounts were rolled out between 2002 and 2005, and thus, beneficiaries had several years of experience with the account when debit cards were first introduced in 2009. Second, both treatment and control beneficiaries were accumulating time with their savings accounts simultaneously, and had accounts for the same amount of time on average. Third, our results from Section IV include account fixed effects, so any time-invariant effect of having the account for a longer period of time would be absorbed.

VII. Conclusion

Debit cards tied to savings accounts could be a promising avenue to facilitate formal savings, as debit cards reduce transaction costs and provide a mechanism to check balances and build trust in financial institutions. We find large effects of debit cards on savings. The debit cards were rolled out over time to beneficiaries of Mexico's cash transfer program Oportunidades. These beneficiaries were already receiving their benefits in a bank account, but for the most part were not saving in their accounts. After two years with a debit card, beneficiaries increased their stock of savings by 2% of annual income. This effect is larger than that of various other savings interventions, including offering commitment devices, no-fee accounts, higher interest rates, lower transaction costs, and financial education.

Both low transaction costs to access savings and trust in banks appear to be necessary but not (individually) sufficient conditions to save in formal financial institutions. Thus, high indirect transaction costs and low trust could potentially explain why a number of studies offering savings accounts with no fees or minimum balance requirements find low take-up and, even among adopters, low use of the accounts. Today, over 100 million poor households receive government cash transfers worldwide, and a growing share is getting their transfers through automatic deposits. In urban areas, these deposits can be withdrawn at ATMs from an extensive ATM infrastructure. Our study suggests that for these populations, the reduction in transaction costs of accessing money and monitoring the bank achieved through debit cards promises to increase financial inclusion and enable the poor to save. We acknowledge that relative to contexts without direct deposits of income into a bank account, with less ATM infrastructure, or with less willingness of retailers to adopt POS terminals, our measured effect of debit cards on savings may be an upper bound.

While our limited time frame prevents us from directly assessing the welfare implications of this policy, a growing literature suggests that enabling the poor to save in formal financial institutions leads to increased welfare through greater investment, wealth accumulation, and ability to cope with shocks, leading to higher long-term consumption. It is worth noting that beneficiaries with debit cards voluntarily used the technology and accumulated savings in their accounts—they could have instead continued withdrawing all of their benefits from the bank branch, as they did prior to receiving the card. This indicates a revealed preference for saving in formal financial institutions once transaction costs are lowered and trust is built. In addition, beneficiary survey responses in the *Trust Survey* indicate that satisfaction with the payment method was higher after receiving the debit card, particularly for those who had the card longer: of beneficiaries with the card for at least 14 months (the median time), 75% indicated that receiving payment by debit card was better than before, and 13% that it was the same as before. In terms of mechanism, our results suggest that cash saved at home was easily spent, potentially due to intra-household bargaining issues. Indeed, after receiving the card, beneficiaries strongly reduced their spending on temptation goods, and the increase in savings was largest in households with low baseline bargaining power for women.

Initial submission: June 18, 2019; Accepted: May 21, 2020
Editors: Stefan Nagel, Philip Bond, Amit Seru, and Wei Xiong

REFERENCES

Anderson, Siwan, and Jean-Marie Baland, 2002, The economics of roscas and intrahousehold resource allocation, *Quarterly Journal of Economics* 117, 963–995.

Ashraf, Nava, 2009, Spousal control and intra-household decision making: An experimental study in the Philippines, *American Economic Review* 99, 1245–1277.

Ashraf, Nava, Diego Aycinena, Claudia Martínez, and Dean Yang, 2015, Savings in transnational households: A field experiment among migrants from El Salvador, *Review of Economics and Statistics* 97, 332–351.

Ashraf, Nava, Dean Karlan, and Wesley Yin, 2006, Deposit collectors, *Advances in Economic Analysis & Policy* 6, 635–672.

Attanasio, Orazio P., Luigi Guiso, and Tullio Jappelli, 2002, The demand for money, financial innovation, and the welfare cost of inflation: An analysis with household data, *Journal of Political Economy* 110, 317–351.

Bachas, Pierre, Paul Gertler, Sean Higgins, and Enrique Seira, 2018, Digital financial services go a long way: Transaction costs and financial inclusion, *American Economic Association Papers & Proceedings* 108, 444–448.

Banerjee, Abhijit, and Sendhil Mullainathan, 2010, The shape of temptation: Implications for the economic lives of the poor, NBER Working Paper 15973.

Baumol, William J., 1952, The transactions demand for cash: An inventory theoretic approach, *Quarterly Journal of Economics* 66, 545–556.

Blumenstock, Joshua, Michael Callen, and Tarek Ghani, 2018, Why do defaults affect behavior? Experimental evidence from Afghanistan, *American Economic Review* 108, 2868–2901.

Borusyak, Kirill, and Xavier Jaravel, 2016, Revisiting event study designs, Harvard University Working Paper.

Breza, Emily, and Arun G. Chandrasekhar, 2019, Social networks, reputation, and commitment: Evidence from a savings monitors experiment, *Econometrica* 87, 175–216.

Bruhn, Miriam, and Inessa Love, 2014, The real impact of improved access to finance: Evidence from Mexico, *Journal of Finance* 69, 1347–1376.

Callen, Michael, Suresh De Mel, Craig McIntosh, and Christopher Woodruff, 2019, What are the headwaters of formal savings? Experimental evidence from Sri Lanka, *Review of Economic Studies* 86, 2491–2529.

Carroll, Christopher D., 1997, Buffer-stock saving and the life cycle/permanent income hypothesis, *Quarterly Journal of Economics* 112, 1–55.

Célérier, Claire, and Adrien Matray, 2019, Bank branch supply, financial inclusion and wealth accumulation, *Review of Financial Studies* 32, 4767–4809.

CGAP, 2012, Social cash transfers and financial inclusion: Evidence from four countries, Consultative Group to Assist the Poor (CGAP) Focus Note 77.

Cole, Shawn, Xavier Giné, Jeremy Tobacman, Peta Topalova, Robert Townsend, and James Vickery, 2013, Barriers to household risk management: Evidence from India, *American Economic Journal: Applied Economics* 5, 104–135.

de Janvry, Alain, Kyle Emerick, Marco Gonzalez-Navarro, and Elisabeth Sadoulet, 2015, Delinking land rights from land use: Certification and migration in Mexico, *American Economic Review* 105, 3125–3149.

Demirgüç-Kunt, Asli, Leora Klapper, Dorothe Singer, Saniya Ansar, and Jake Hess, 2018, The Global Findex Database 2017: Measuring financial inclusion and the fintech revolution, World Bank Report.

Drexler, Alejandro, Greg Fischer, and Antoinette Schoar, 2014, Keeping it simple: Financial literacy and rules of thumb, *American Economic Journal: Applied Economics* 6, 1–31.

Dupas, Pascaline, Sarah Green, Anthony Keats, and Jonathan Robinson, 2016, Challenges in banking the rural poor: Evidence from Kenya's Western Province, in Sebastian Edwards, Simon Johnson, and David N. Weil, eds.: Modernization and Development (University of Chicago Press, Chicago).

Dupas, Pascaline, Dean Karlan, Jonathan Robinson, and Diego Ubfal, 2018, Banking the unbanked? Evidence from three countries, *American Economic Journal: Applied Economics* 10, 257–297.

Dupas, Pascaline, and Jonathan Robinson, 2013, Savings constraints and microenterprise development: Evidence from a field experiment in Kenya, *American Economic Journal: Applied Economics* 5, 163–192.

FDIC, 2016, Bank efforts to serve unbanked and underbanked consumers. https://www.fdic.gov/consumers/community/research/qualitativeresearch_may2016.pdf

Federal Reserve, 2017, Report on the economic well-being of U.S. households in 2016. https://www.federalreserve.gov/publications/files/2016-report-economic-well-being-us-households-201705.pdf

Galiani, Sebastian, Paul Gertler, and Ernesto Schargrodsky, 2005, Water for life: The impact of the privatization of water services on child mortality, *Journal of Political Economy* 113, 83–120.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales, 2004, The role of social capital in financial development, *American Economic Review* 49, 526–556.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales, 2008, Trusting the stock market, *Journal of Finance* 63, 2557–2600.

Haushofer, Johannes, and Jeremy Shapiro, 2016, The short-term impact of unconditional cash transfers to the poor: Experimental evidence from Kenya, *Quarterly Journal of Economics* 131, 1973–2042.

Higgins, Sean, 2019, Financial technology adoption, Working Paper.

Jack, William, and Tavneet Suri, 2014, Risk sharing and transactions costs: Evidence from Kenya's mobile money revolution, *American Economic Review* 104, 183–223.

Jakiela, Pamela, and Owen Ozier, 2016, Does Africa need a rotten kin theorem? Experimental evidence from village economies, *Review of Economic Studies* 83, 231–268.

Johnson, Eric J., Stephan Meier, and Olivier Toubia, 2019, What's the catch? Suspicion of bank motives and sluggish refinancing, *Review of Financial Studies* 32, 467–495.

Karlan, Dean, Jake Kendall, Rebecca Mann, Rohini Pande, Tavneet Suri, and Jonathan Zinman, 2016, Research and impacts of digital financial services, NBER Working Paper 22633.

Karlan, Dean, Markus Mobius, Tanya Rosenblat, and Adam Szeidl, 2009, Trust and social collateral, *Quarterly Journal of Economics* 124, 1307–1361.

Karlan, Dean, Aishwarya Lakshmi Ratan, and Jonathan Zinman, 2014, Savings by and for the poor: A research review and agenda, *Review of Income and Wealth* 60, 36–78.

Karlan, Dean, Beniamino Savonitto, Bram Thuysbaert, and Christopher Udry, 2017, Impact of savings groups on the lives of the poor, *Proceedings of the National Academy of Sciences* 114, 3079–3084.

Karlan, Dean, and Jonathan Zinman, 2018, Price and control elasticities of demand for savings, *Journal of Development Economics* 130, 145–159.

Kast, Felipe, Stephan Meier, and Dina Pomeranz, 2018, Saving more in groups: Field experimental evidence from Chile, *Journal of Development Economics* 133, 275–294.

Kast, Felipe, and Dina Pomeranz, 2014, Saving more to borrow less: Experimental evidence from access to formal savings accounts in Chile, HBS Working Paper 14-001.

Mehrotra, Rahul, Vincent Somville, and Lore Vandewalle, 2021, Increasing trust in bankers to enhance savings: Experimental evidence from India, *Economic Development and Cultural Change* 69, 623–644.

Morduch, Jonathan, and Rachel Schneider, 2017, The Financial Diaries: How American Families Cope in a World of Uncertainty (Princeton University Press, Princeton, NJ).

Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar, 2016, Building state capacity: Evidence from biometric smartcards in India, *American Economic Review* 106, 2895–2929.

Olafsson, Arna, and Michaela Pagel, 2017, The ostrich in us: Selective attention to financial accounts, income, spending, and liquidity, Working Paper 23945, National Bureau of Economic Research.

Osili, Una Okonkwo, and Anna Paulson, 2014, Crises and confidence: Systemic banking crises and depositor behavior, *Journal of Financial Economics* 111, 646–660.

Parker, Susan W., and Petra E. Todd, 2017, Conditional cash transfers: The case of Progresa/Oportunidades, *Journal of Economic Literature* 55, 866–915.

Prina, Silvia, 2015, Banking the poor via savings accounts: Evidence from a field experiment, *Journal of Development Economics* 115, 16–31.

Schaner, Simone, 2015, Do opposites detract? Intrahousehold preference heterogeneity and inefficient strategic savings, *American Economic Journal: Applied Economics* 7, 135–174.

Schaner, Simone, 2017, The cost of convenience? Transaction costs, bargaining, and savings account use in Kenya, *Journal of Human Resources* 52, 919–943.

Schaner, Simone, 2018, The persistent power of behavioral change: Long-run impacts of temporary savings subsidies for the poor, *American Economic Journal: Applied Economics* 10, 67–100.

Seshan, Ganesh, and Dean Yang, 2014, Motivating migrants: A field experiment on financial decision-making in transnational households, *Journal of Development Economics* 108, 119–127.

Somville, Vincent, and Lore Vandewalle, 2018, Saving by default: Evidence from a field experiment in rural India, *American Economic Journal: Applied Economics* 10, 39–66.

Stein, Luke C.D., and Constantine Yannelis, 2020, Financial inclusion, human capital, and wealth accumulation: Evidence from the Freedman's Savings Bank, *Review of Financial Studies* 33, 5333–5377.

Suri, Tavneet, and William Jack, 2016, The long-run poverty and gender impacts of mobile money, *Science* 354, 1288–1292.

Tobin, James, 1956, The interest-elasticity of transactions demand for cash, *Review of Economics and Statistics* 38, 241–247.

Zinman, Jonathan, 2009, Credit or debit? *Journal of Banking and Finance* 33, 358–366.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix S1: Internet Appendix.
Replication Code.