NBER WORKING PAPER SERIES

BANKING ON TRUST: HOW DEBIT CARDS ENABLE THE POOR TO SAVE MORE

Pierre Bachas Paul Gertler Sean Higgins Enrique Seira

Working Paper 23252 http://www.nber.org/papers/w23252

NATIONAL BUREAU OF ECONOMIC RESEARCH 1050 Massachusetts Avenue Cambridge, MA 02138 March 2017

This study would not have been possible without the generous time, access to data, and tireless answers to questions provided by officials in Mexico's government bank Bansefi and the conditional cash transfer program Prospera (formerly Oportunidades). At Bansefi, we are indebted to 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 Bibek Adhikari, Jesse Anttila-Hughes, Stefano Barbieri, Alan Barreca, Richard Blundell, Pablo Celhay, Carlos Chiapa, Laura Chioda, Natalie Cox, Alain de Janvry, John Edwards, Gerardo Escaroz, Fred Finan, Keith Finlay, Louise Fox, Jess Goldberg, Emilio Gutierrez, Jens Hainmueller, Anders Jensen, Anne Karing, Dean Karlan, Supreet Kaur, Matt Larsen, John Loeser, Paty López, Nora Lustig, Jeremy Magruder, Justin McCrary, Shaun McRae, Atif Mian, Ted Miguel, John Morgan, Doug Nelson, Christine Parlour, Gérard Roland, Whitney Ruble, Betty Sadoulet, Frank Schilbach, Todd Schoellman, Steven Sheffrin, Isidro Soloaga, Ben Sperisen, Chele Teruel, Bruce Wydick, Stephen Younger, Jonathan Zinman, and numerous seminar participants. We are also grateful to Ignacio Camacho, Ernesto Castillo, Oscar Cuellar, Bernardo García, Austin Farmer, Isaac Meza, Mel Reitcheck, and especially Joel Ferguson 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). All authors declare that they have no relevant or material financial interests that relate to the research described in this paper. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peerreviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Pierre Bachas, Paul Gertler, Sean Higgins, and Enrique Seira. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Banking on Trust: How Debit Cards Enable the Poor to Save More Pierre Bachas, Paul Gertler, Sean Higgins, and Enrique Seira NBER Working Paper No. 23252 March 2017 JEL No. D14,D83,G21,O16

ABSTRACT

Trust is an essential element of economic transactions, but trust in financial institutions is low, especially among the poor. Debit cards provide not only easier access to savings, but also a mechanism to monitor bank account balances and thereby build trust in a financial institution. We study a natural experiment in which debit cards are rolled out to beneficiaries of a Mexican conditional cash transfer program whose benefits are already directly deposited into a savings account. Using administrative data on over 340,000 bank accounts over four years, we find that prior to receiving a debit card, beneficiaries do not save in these accounts. Beneficiaries then begin to increase their savings after 9 to 12 months with the card. During this initial stagnant period, they use the card to check their balances frequently, and the number of checks decreases over time as their reported trust in the bank increases. After 1 to 2 years, the debit card causes the savings rate to increase by 3 to 5 percent of income. Using household survey panel data, we find that this effect represents an increase in overall savings.

Pierre Bachas Princeton Economics Department Office 214 Princeton, NJ 08542 pbachas@princeton.edu

Paul Gertler Haas School of Business University of California, Berkeley Berkeley, CA 94720 and NBER gertler@haas.berkeley.edu Sean Higgins Center for Effective Global Action University of California, Berkeley Berkeley, CA 94720 seanhiggins@berkeley.edu

Enrique Seira Centro de Investigación Económica ITAM Ave. Santa Teresa # 930 Mexico, D. F. 10700, Mexico enrique.seira@gmail.com Virtually every commercial transaction has within itself an element of trust.... It can be plausibly argued that much of the economic backwardness in the world can be explained by the lack of mutual confidence.

-Kenneth Arrow (1972)

1 Introduction

Trust is an essential element of economic transactions and an important driver of economic development (Knack and Keefer, 1997; La Porta et al., 1997; Algan and Cahuc, 2010). It is particularly crucial in financial transactions where people pay money in exchange for promises, and essential where the legal institutions that enforce contracts are weak (McMillan and Woodruff, 1999; Karlan et al., 2009). Given the nature of financial decisions, it is not surprising that trust has been shown to be key to stock market participation (Guiso, Sapienza, and Zingales, 2008), use of checks instead of cash (Guiso, Sapienza, and Zingales, 2004), and decisions to not withdraw deposits from financial institutions in times of financial crisis (Iyer and Puri, 2012; Sapienza and Zingales, 2012).

Trust in financial institutions, meanwhile, is low. Majorities in close to half of the countries included in the World Values Survey report lack of confidence in banks. Trust is especially low among the poor: in Mexico, the location of our study, 71% of those with less than a primary school education report low trust in banks, compared to 55% of those who completed primary school and 46% of those who completed university (Figure I).

Lack of trust in financial institutions may not be unfounded. Cohn, Fehr, and Maréchal (2014) provide evidence that the banking industry fosters a culture of dishonesty relative to other industries. In Mexico, bankers loot money by directing lending to "related parties," i.e. bank shareholders and their firms (La Porta, López-de-Silanes, and Zamarripa, 2003). Mexican newspapers report many instances of outright bank fraud where depositors have lost their savings. For example, an extensively covered scandal involved Ficrea, whose majority shareholder reportedly stole US\$ 200 million from savers (CNBV, 2014). Bank fraud is frequently reported in the press, with at least 275 news stories about 32 unique events of savings fraud published in 2014 and 2015 alone.¹ Tellingly, articles that provide financial advice in Mexican newspapers have titles like "How to Save for Your

¹We scraped the online news archives of all electronic newspapers and news websites in Mexico using several keywords, and then filtered the results by hand to keep only relevant stories. The scraping resulted in 1392 stories in 121 newspapers from 2014-2015 that matched our keywords, of which 275 stories from 35 newspapers directly reported on bank fraud.

Graduation and Avoid Fraud" and "Retirement Savings Accounts, with Minimal Risk of Fraud."

When fraud is rampant and contract enforcement poor, trust plays an even larger role (Guiso, Sapienza, and Zingales, 2004; Karlan et al., 2009) and people are understandably reluctant to use financial institutions (Bohnet, Herrmann, and Zeckhauser, 2010). At the country level, trust is strongly associated with the proportion of the population that do save in formal bank accounts (Figure II). Along with fees and minimum balance requirements, trust is frequently listed by the poor as a primary reason for not saving in formal bank accounts (e.g., Dupas et al., 2016a). Lack of trust could also explain why randomized field experiments in three countries have found that even among people who take up accessible and free formal savings products, account use is low (Dupas et al., 2016b).² Despite its importance, finding ways to improve trust in financial institutions has not been extensively studied (Karlan, Ratan, and Zinman, 2014).

While trust is important, it is not an innate characteristic but rather can be influenced through experience and information (Hirschman, 1984; Williamson, 1993; Attanasio, Pellerano, and Reyes, 2009). Debit cards (and mobile money) provide a low-cost technology to monitor account balances and thereby build trust that a bank is not explicitly stealing deposits or charging unexpectedly large hidden fees.³ We hypothesize that new debit card clients first use the cards to check balances and thereby establish trust, after which they take advantage of the cards' lower transaction costs to use the services of formal financial institutions. In this sense, we argue that building trust in a financial institution is a necessary condition for the use of formal financial services; i.e., financial inclusion requires trust.

We examine this hypothesis in the context of a natural experiment in which debit cards tied to savings accounts were rolled out geographically over time to beneficiaries of the Mexican conditional cash transfer program Oportunidades. The phased geographic rollout provides plausibly exogenous variation in assignment of debit cards to beneficiaries in a difference-in-differences context. Before the rollout, beneficiaries had been receiving their transfers through savings accounts without debit cards, and very rarely used their accounts to save. Instead, they typically withdrew almost the full amount of the transfer shortly after receiving it. This is consistent with findings from other

²Trust is hypothesized as one channel through which no-fee accounts led to increased saving in Prina (2015).

³Previous studies on debit cards and mobile money in developing countries have focused on the effect of the lower transaction costs facilitated by these technologies to make purchases (Zinman, 2009), access savings and remittances (Suri, Jack, and Stoker, 2012; Schaner, forthcoming), and transfer money (Jack, Ray, and Suri, 2013; Jack and Suri, 2014), but not their capacity to monitor and build trust in financial institutions.

countries such as Brazil, Colombia, India, Niger, and South Africa, in which cash transfers are also paid through bank or mobile money accounts and recipients generally withdraw the entire transfer amount in one lump sum withdrawal each pay period (Bold, Porteous, and Rotman, 2012; Aker et al., 2016; Muralidharan, Niehaus, and Sukhtankar, 2016).

This paper makes four contributions. First, we show that debit cards cause a large and significant increase in savings in formal financial institutions: after a delay, beneficiaries with debit cards save 3–5% more of their income each period. Second, we find that this increase in savings is driven in large part by clients using the debit card to first monitor account balances and thereby build trust that their money is safe. Once trust is established, they take advantage of the reduced transaction costs associated with debit cards and increase the amount of money held in their bank accounts.⁴ Third, we find that the observed higher savings in the bank constitute an increase in total savings and not just a substitution from other savings vehicles. Finally, our study uses a much larger sample than most of the literature, with broad geographic coverage across the country.

The size of the effect we observe is larger than that of other savings interventions studied in the literature, including offering commitment devices, no-fee accounts, higher interest rates, lower transaction costs, and financial education. The one exception is Suri and Jack (2016), who study the impact of mobile money—another technology that enables clients to more easily check account balances and build trust. They find that female-headed households increase their savings rate by 3% of income on average, after six years of mobile money exposure.

For the analysis, we use high frequency administrative data on bank transactions for over 340,000 beneficiary accounts in 357 bank branches nationwide over 4 years, as well as several surveys of beneficiaries. We compare three groups of beneficiaries based on the rollout of debit cards: beneficiaries in two treatment waves receive debit cards one year apart, while those in the control group receive a debit card at the end of our study period. Using the administrative data, we find that beneficiaries initially use debit cards to check account balances without increasing their savings. Over time, the frequency of account balance checks falls, and after a delay of 9 to 12 months in wave 1, savings rates rise. We observe a faster effect in wave 2, and find evidence that this faster effect is due to

⁴In our context, debit cards reduce the *indirect* transaction costs of accessing money in the bank account, as savings can be withdrawn at any bank's ATM rather than only at government bank branches, which are often far from beneficiaries. In contrast, Schaner (forthcoming) provides ATM cards that reduce *direct* transaction costs: higher withdrawal fees are charged by bank tellers in her study, and the only ATMs at which the cards can be used are located at bank branches of the corresponding bank.

information spillovers, as wave 2 beneficiaries living farther from wave 1 localities still have a delay before they begin saving. We find that after 1–2 years with the card, the share of total income saved each payment period increases by 3–5 percentage points. These savings rates closely align with the savings goals individuals set in Breza and Chandrasekhar (2015).

The delayed initiation of saving suggests some kind of learning. We use beneficiary survey data to explore three kinds of learning: (i) learning to trust the bank, (ii) learning to use the ATMs, their location, and associated transaction costs, and (iii) learning that the program will not drop beneficiaries who accumulate savings. We find support for the "learning to trust" hypothesis. Specifically, beneficiaries who have had their debit cards for less time report significantly lower rates of trusting the bank than beneficiaries who have had their debit cards longer. On the other hand, we find no support for the other forms of learning.

To establish a direct link between trust and increased saving, we merge the administrative data on account balances and transactions with the beneficiary survey reporting trust in the bank. Since trust is both endogenous to the savings decision and susceptible to measurement error, we instrument trust with a set of dummies for timing of debit card receipt. We find that beneficiaries who are induced to trust the bank as a result of having the card longer save an additional 3% of their income. To our knowledge, this provides the first direct causal estimate in the literature of the effect of trust in financial institutions on formal savings.

We then test whether the increase in bank account balances is an increase in total savings or a substitution from other forms of saving, both formal and informal. Using household survey panel data, we find that after close to one year beneficiaries in the treatment group increase their savings rate by 5% of income relative to the control group, which is very close in magnitude to the effect we see in the administrative account data. We find no differential change in income or assets in the treatment group compared to the control, but rather that the increase in savings is financed though reduced current consumption. Hence, the increase in formal bank account savings does not appear to crowd out other forms of saving (consistent with results in Dupas and Robinson, 2013a; Ashraf, Karlan, and Yin, 2015; Kast, Meier, and Pomeranz, 2016).

Given our results, government cash transfer programs could be a promising channel to increase financial inclusion and enable the poor to save, not only because of the sheer number of the poor 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 accounts, debit cards, and mobile money (Aker et al., 2016; Haushofer and Shapiro, 2016; Muralidharan, Niehaus, and Sukhtankar, 2016). Furthermore, debit cards combined with ATMs or point-of-sale terminals and mobile phones combined with mobile money platforms are low-cost technologies that can be used to check balances and build trust in financial institutions. These technologies are simple, prevalent, and potentially scalable to millions of cash transfer recipients worldwide.

2 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, forthcoming). The program provides bimonthly cash transfers to poor families conditional on sending their children to school and having preventive health check-ups. The program seeks to alleviate poverty in the short term and break the intergenerational transmission of poverty by encouraging families to invest in the human capital of the next generation. It began in rural Mexico in 1997 under the name Progresa, and later expanded to urban areas starting in 2002. Today, nearly one-fourth of Mexican households receive benefits from Oportunidades (Levy and Schady, 2013).⁵

As it expanded to urban areas in 2002–2005, Oportunidades opened savings accounts in banks for beneficiaries in a portion of urban localities, and began depositing the transfers directly into those accounts. The original motives for paying through bank accounts were to (i) decrease corruption as automatic payments through banks lower the ability of local officials to skim off benefits⁶ and of local politicians to associate themselves with the program through face-to-face contact with recipients when they receive their transfers, (ii) decrease long wait times for recipients who previously had to show up to a "payment table" on a particular day to receive their benefits, and (iii) decrease robberies and assaults of program officers and recipients transporting cash on known days.

By the beginning of 2005, over one million families received their benefits directly deposited into savings accounts in Bansefi, a government bank created to increase savings and financial inclusion

 $^{^{5}}$ Oportunidades was recently rebranded as Prospera. We use the name that was in place during our study period.

⁶Consistent with this concern, Muralidharan, Niehaus, and Sukhtankar (2016) find that paying government cash transfers through biometric "smartcards" in India led to a 40% reduction in program leakages to corrupt officials.

among underserved populations (Figure III).⁷ The Bansefi savings accounts have no minimum balance requirement or monthly fees and pay essentially no interest.⁸ Before the introduction of debit cards, beneficiaries could only access their money at Bansefi bank branches. Because there are only about 500 Bansefi branches nationwide and many beneficiaries live far from their nearest branch, accessing their accounts involved large transaction costs for many beneficiaries. Overall, the savings accounts were barely used prior to the introduction of debit cards. Prior to the rollout of debit cards, the average number of deposits per bimester⁹ is almost exactly one—including the deposit from Oportunidades—and over 90% of clients make one withdrawal each bimester, withdrawing on average 99% of the transfer.

In 2009, the government announced that they would issue Visa debit cards to beneficiaries who were receiving their benefits directly deposited into Bansefi savings accounts. The cards enable account holders to withdraw cash and check account balances at any bank's ATM, as well as make electronic payments at any store accepting Visa. Beneficiaries can make two free ATM withdrawals per bimester at any bank's ATM; additional ATM withdrawals are charged a fee that varies by bank. When Bansefi distributed the debit cards, they also provided beneficiaries with a training session on how and where to use the cards.¹⁰ The training session did not vary over time and did not encourage recipients to save.

In 275 out of Mexico's 550 urban localities, beneficiaries received their benefits in bank accounts prior to the rollout of debit cards. Debit cards tied to these bank accounts were rolled out to approximately 75,000 beneficiaries in 143 localities in 2009 (wave 1) and to an additional 170,000 beneficiaries in 88 localities in late 2010 (wave 2). Another 100,000 beneficiaries in the remaining localities received cards between November 2011 and February 2012 (control group), immediately after the end date of our study period. The map in Figure B.1 shows that the treatment and control waves had substantial geographical breadth.

The sequence with which beneficiaries in localities switched to debit cards was determined as a function of the proportion of households in the locality that were eligible for the program but were

⁷Originally Oportunidades partnered with two banks: Bansefi, a government bank, and Bancomer, a commercial bank. However, working with a commercial bank proved to be difficult, and Oportunidades phased out the Bancomer accounts and transferred them to Bansefi by mid-2006.

 $^{^{8}}$ Nominal interest rates were between 0.09 and 0.16% per year compared to an inflation rate of around 5% per year during our sample period.

⁹The program is paid in two-month intervals, which we refer to throughout the paper as bimesters. The Spanish word *bimestre* is more common than its English cognate, and is used by Bansefi and Oportunidades.

¹⁰See Appendix A for a sample of the materials that beneficiaries received together with their cards.

not yet receiving benefits. This is because the introduction of debit cards to existing recipients was coupled with an effort to incorporate more beneficiaries. Table I compares the means of localitylevel variables and account-level variables from the control, wave 1, and wave 2 localities using data from (i) the population census from 2005, (ii) poverty estimates from Oportunidades based on the same census, (iii) Bansefi branch locations from 2008, and (iv) the administrative account data on average balances and transactions from Bansefi in 2008. Column 6 shows the p-value of an F-test of equality of means. Because the rollout was not random, it is not surprising that there are some differences across treatment and control localities: treatment localities are slightly larger and beneficiaries in these localities receive higher transfer amounts. The share of the transfer withdrawn is high, ranging from 99.2 to 99.4% of the transfer, indicating very low savings in the account prior to receiving the card.

3 Data Sources

We use a rich combination of administrative and household survey data sources. To examine the effect of debit cards on savings and account use, we analyze account-level transaction data from Bansefi for 343,204 accounts at 380 Bansefi branches over a four-year period, from November 2007 to October 2011. These data include the monthly average savings balance, the date and amount of each transaction made in the account (including Oportunidades transfers), the date the account was opened, and the month the card was awarded to the account holder. Note that the dates that the account was opened and the debit card obtained are determined exogenously by Oportunidades, not endogenously by the beneficiary; the average account had been opened 5.3 years before receiving a card. Figure IIIa shows the timing of the administrative Bansefi account balance and transaction data relative to the rollout of debit cards.

To test whether the delayed savings effect can be explained by learning to trust the bank, learning to use the technology, learning the program rules, or other types of learning, we use a combination of the administrative transaction data from Bansefi and two household surveys conducted by Oportunidades. Specifically, we first use the administrative transaction data to examine balance check behavior and test for a within-account correlation between balance checks and saving. We then use the Payment Method Survey, a household survey conducted by Oportunidades in 2012 aimed at eliciting information about beneficiaries' satisfaction with and use of the debit cards, to examine mechanisms behind and outcomes of the different types of learning that could be taking place. This survey asks beneficiaries to report the number of balance checks they make each period, whether they find it hard to use the ATM, get help using the ATM, and know their PIN, and the fees they are charged for balance checks and withdrawals. Finally, we use the Survey of Urban Households' Sociodemographic Characteristics (ENCASDU), conducted by Oportunidades in late 2010, to investigate beneficiaries' self-reported reasons for not saving in their Bansefi accounts. We also merge the ENCASDU with administrative Bansefi data at the beneficiary/account level to study the direct relationship between self-reported trust in the bank and actual savings levels.

Finally, to explore whether the increased savings in the Bansefi accounts is an increase in overall savings or a substitution from other forms of saving, we use the Survey of Urban Household Characteristics (ENCELURB), a panel survey with three pre-treatment waves in 2002, 2003, and 2004, and one post-treatment wave conducted from November 2009 to 2010. This survey has comprehensive modules on consumption, income, and assets. We merge these data with administrative data from Oportunidades on beneficiary status and the dates that debit cards were distributed in each locality. Figure IIIb shows the timing of the household survey data relative to the rollout of debit cards.

4 Effect of Debit Cards on Savings Behavior and Account Use

In this section, we use the administrative data from Bansefi on average monthly balances and all transactions in 343,204 accounts of Oportunidades beneficiaries to estimate the dynamic effect of debit cards on accumulated savings in these formal bank accounts, on use of the accounts through transactions (deposits and withdrawals), and on the savings rate. We exploit the phased rollout of debit cards to identify the causal effect of debit cards on savings and account use in a difference-in-differences framework.

4.1 The Stock of Savings (Account Balances)

The raw data on average account balances display the fundamental patterns of interest. Because Oportunidades payments are made every two months, in Figure IV we plot the bimonthly average account balances of beneficiaries in the treatment and control groups over time.¹¹ Panel (a) presents

¹¹These are obtained using the average monthly account balances provided by Oportunidades and averaging within accounts over adjacent months to obtain bimonthly averages.

the comparison between wave 1 and the control, and panel (b) between wave 2 and the control. The dashed vertical lines indicate the time at which beneficiaries in treatment localities received debit cards. For both waves, average balances are almost identical in both levels and trends prior to the introduction of debit cards. Strikingly, treatment group savings in wave 1 rise dramatically relative to the control group after an initial period of about 10 months with the debit card. After two years with the card, wave 1 beneficiaries have average balances that are about four times larger than those of the control group. We see a similar pattern in wave 2, but due to the later switch to cards in wave 2, we have fewer periods after treatment to observe the dynamic effect over time. We also observe an effect of the debit cards on savings more quickly in wave 2, which we explore further in Section 4.2.

In order to estimate the causal effect of the debit card on account balances, we estimate a difference-in-differences specification that allows the treatment effect to vary over time. We control for common macro shocks by including time fixed effects, and for time-invariant individual heterogeneity with individual account fixed effects. Specifically, we estimate

$$Balance_{it} = \lambda_i + \delta_t + \sum_k \phi_k T_{j(i)} \times \mathbb{I}(t=k) + \varepsilon_{it}$$
(1)

separately for wave 1 and wave 2 (each compared to the control group), where $Balance_{it}$ is the average balance in account *i* over period *t*, the λ_i are account-level (i.e., beneficiary) fixed effects, the δ_t are time-period fixed effects, $T_{j(i)} = 1$ if locality *j* in which account holder *i* lives is a treatment (i.e., wave 1 or wave 2) locality, and $\mathbb{I}(t = k)$ are time period dummies.¹² Thus, the $T_{j(i)} \times \mathbb{I}(t = k)$ terms pick up the difference in balances between treatment and control localities in each period.

We average the data within four-month periods because not all beneficiaries receive their payments in each calendar bimester: some payments are shifted to the latter part of the prior bimester in some localities, resulting in some bimesters with double payments and others with no payments.¹³ Because we have four years of data, this leaves us with 12 four-month periods. Since one time period dummy and one $T_{j(i)} \times \mathbb{I}(t = k)$ term must be omitted from (1), we follow the standard procedure

¹²Following other papers measuring savings (e.g., de Mel, McIntosh, and Woodruff, 2013; Chetty et al., 2014; Dupas et al., 2016a; Kast, Meier, and Pomeranz, 2016; Karlan and Zinman, 2016), we winsorize average balances to avoid results driven by outliers; our main results winsorize at the 95th percentile, and the results are robust to other cut-offs. To avoid truncating a true treatment effect or time trend, we winsorize within each wave and time period.

¹³This payment shifting happens for various reasons, including for local, state, and federal elections, as a law prohibits Oportunidades from distributing cash transfers during election periods to prevent corruption.

of omitting the period immediately preceding the change to cards. We estimate cluster-robust standard errors, clustering ε_{it} by Bansefi branch. The coefficients of interest are the ϕ_k , which measure the average difference in balances between the treatment and control group in period k.

Figure V plots the ϕ_k coefficients and their 95% confidence intervals. As we saw in Figure IV, it is clear that the pre-treatment levels and trends in account balances are not different between treatment and control groups prior to the debit card rollout. This is important, as the identifying assumption in our difference-in-differences model is that the beneficiaries who receive the debit card would have had the same trend in average balances as the control group in the absence of treatment. While this assumption is inherently untestable, the fact that the levels and trends of the dependent variable are very similar before treatment makes a strong case for the parallel trends assumption. Having two waves is also helpful for identification as it provides a test of validity for other time periods and populations and suggests that results are not due to a one-time macro shock.

In wave 1, there is no difference in average account balances between the treatment and control groups for about 8 months after receiving debit cards, beyond which balances start to increase significantly. In Section 5 we test three different types of learning that could explain the delayed effect, and find evidence that beneficiaries are using the debit cards to monitor and build trust in the bank; after building trust, they take advantage of the cards' lower transaction costs (due to being able to withdraw at any bank's ATM rather than the nearest Bansefi branch) and use the account to save. After nearly two years with the card, balances are about 1,400 pesos higher for the treatment group than for the control. Furthermore, dynamic savings effects appear to be at play: once beneficiaries begin to save, their balances increase at a decreasing rate. We explore these dynamics further in Section 4.4. In wave 2, savings begin to rise much more quickly after receiving the card; in Section 4.2 we show that this is likely due to information spillovers from wave 1 beneficiaries.

4.2 Information Spillovers

In this section, we test whether the faster effect of debit cards on the accumulation of savings in wave 2 compared to wave 1 can be explained by information spillovers: i.e. as beneficiaries in wave 1 learned about the bank's trustworthiness and the value of savings, they may have shared this information with family members and friends living in nearby wave 2 (and control) localities.¹⁴ Thus, wave 2 beneficiaries may have already had higher levels of trust when they received debit cards than did wave 1 beneficiaries; if trust is a necessary condition for saving but not sufficient when transaction costs are too high, we would then see wave 2 beneficiaries start saving once the debit card lowered their transaction costs. Control beneficiaries would be less likely to save in response to information spillovers, as their transaction costs remain high absent a debit card.

To test this, we split wave 2 and control beneficiaries into two groups: those who are closer to and farther from wave 1 localities. We then test whether the effect of debit cards is faster and larger for wave 2 beneficiaries who are close to wave 1 localities (relative to close control beneficiaries) than for wave 2 beneficiaries who are far from wave 1 localities (relative to far control beneficiaries). Specifically, we estimate (1) separately for wave 2 and control beneficiaries who are close to wave 1 localities, and for wave 2 and control beneficiaries who are far from wave 1 localities.

Because we do not know *ex ante* the distance that information about trustworthiness might travel, we use various percentile cut-offs to separate wave 2 and control beneficiaries that are closer vs. farther from wave 1 localities. Specifically, we test different percentile cut-offs between 10% and 90% of the distribution of road distances to the centroid of wave 1 localities. We find that, when we split the population at the 80th percentile or higher, the effect of debit cards on the stock of savings is faster and larger for wave 2 beneficiaries who are close to wave 1 localities than for those who are farther. Figure VI presents the results using the 80th percentile (which corresponds to a road distance of 154 kilometers) as the cut-off.

In all three periods after receiving cards, the effect of the card on savings for close beneficiaries is larger than that for far beneficiaries, and the difference between the effect sizes is statistically significant. Furthermore, there is a 4-month delay for far beneficiaries before the card has a savings effect, and no delay for close beneficiaries. Given that our tests at different cut-offs reveal a faster effect for the closer 80% of wave 2 and control beneficiaries, it is not surprising that overall, we see an immediate effect of the debit card on savings in wave 2. Taken together, these results suggest that the faster effect of debit cards in wave 2 can in part be explained by information spillovers: after wave 1 beneficiaries learn, they share this information with beneficiaries in nearby localities.

¹⁴The faster effect is not explained by changes in the information given to beneficiaries during the rollout: we obtained detailed documents on the information given to beneficiaries both when the cards were distributed and during bimonthly chats conducted by Oportunidades, and confirmed that there were no substantial changes in this material over time.

4.3 Transactions

By lowering indirect transaction costs, debit cards should lead to more transactions, as predicted by theory (Baumol, 1952; Tobin, 1956) and past empirical evidence (Attanasio, Guiso, and Jappelli, 2002; Alvarez and Lippi, 2009; Schaner, forthcoming). This is indeed what we find. Figure VII presents the distribution of the number of withdrawals per bimester, before and after receiving the card. Prior to receiving the card, over 90% of beneficiaries made a single withdrawal per bimester. After receiving the card, 65% of beneficiaries continue to make just one withdrawal, but 27% make 2 withdrawals, 6% make 3 withdrawals, and 2% make 4 or more withdrawals.¹⁵ Meanwhile, the distribution of the number of withdrawals in the control group does not change over time.

The effect on withdrawals is *immediate*, as would be expected from the instantaneous change in transaction costs induced by the card. Figure VIII shows the average number of withdrawals per Oportunidades deposit in each bimester.¹⁶ Prior to receiving the card, beneficiaries in both the treatment and control groups average close to 1 withdrawal per bimester. Immediately after receiving the card, this figure jumps to an average of about 1.4 withdrawals the period after receiving the card, then remains relatively constant between 1.3 and 1.4 withdrawals in most periods; in the control group, it remains constant at about 1 withdrawal per bimester.

On the contrary, there is no effect on client deposits: Figure IX shows that 99% of accounts have zero client deposits per bimester before and after receiving the card. Account holders thus do not add savings from other sources of income to their Bansefi accounts. This finding is not surprising, since beneficiaries receive one-fifth of their income from the Oportunidades program on average, so unless the optimal savings rate in a particular period is higher than 20% of income, there is no reason to deposit savings from other income sources in the account.

The fact that debit cards increase the number of withdrawals over a period could lead to a mechanical increase in the average balance within a payment period without increasing savings over time between periods. Suppose, for example, that an individual begins a period with a balance of

¹⁵These percentages average over waves 1 and wave 2, combining all account-bimester pairs after the account receives a debit card. Figure VII shows the effects separately for the two waves, again averaging within each wave over all bimesters that are either before or after receiving the card. Recall that the first two withdrawals per bimester are free at any bank's ATM, but subsequent withdrawals are charged a fee, which may explain why so few beneficiaries make more than two withdrawals even after receiving the card.

¹⁶Because of the shifting of some payments to other bimesters described earlier, we graph the number of withdrawals *per Oportunidades deposit* rather than the absolute number of withdrawals. The latter figure looks similar, but—as expected—with higher numbers of withdrawals in bimesters in which this payment shifting leads to more than one deposit.

0, receives an Oportunidades deposit during the period, and withdraws the full amount on the day the funds are deposited. In this case, the average balance over the period is zero. Compare this to an individual who withdraws half the money the day it is deposited and the other half in the middle of the period. In this case, the average balance would equal one-quarter of the transfer amount (since half of the transfer was left in the account for half of the period). In both cases, however, there is no increase in overall savings, if savings are defined as the balance carried over from one period to the next. The size of this mechanical effect depends on the number, timing, and amounts of the withdrawals. Using data on the timing and amount of each transfer and withdrawal from the administrative transactions data, we calculate this mechanical effect for each account-period pair (see Appendix C). We then subtract out the mechanical effect from the average balance to obtain a better measure of the stock of savings, which we call "net balance." The results for (1) using net balances—which we denote Net $Balance_{it}$ —show the same pattern over time; after nearly two years with the card, net balances in wave 1 accounts are approximately 1000 pesos higher than in the control group (Figure B.2). The robustness of results to using net balances, as well as the finding that withdrawal patterns change immediately after receiving the card then remain fairly constant over time, confirm that the change in savings over time is not mechanically driven by changes in the pattern of transactions. For the remainder of the paper, we use the net balances measure.

4.4 Savings Rates

In this section, we examine the impact of debit cards on the savings rate—i.e., the flow of savings as a share of income. There are a number of reasons why households save, including to smooth consumption over the life cycle (Modigliani, 1986), accumulate money for non-divisible purchases of durables in the face of credit constraints (Rosenzweig and Wolpin, 1993), and build a precautionary buffer stock to insure consumption against unexpected shocks (Deaton, 1991). While there is little evidence that life-cycle saving is an important generator of wealth in developing countries, credit constraints make precautionary saving and saving to purchase durables particularly important (Deaton, 1992; Rosenzweig, 2001).¹⁷ The key insight for our purpose is that both the precautionary saving and saving to purchase durables motives lead to a savings target, and as a result, an individual's savings rate is decreasing in her stock of savings as it approaches the target (Carroll, 1997;

¹⁷Even in rich countries, Skinner (1988) finds that precautionary savings constitute a large share of overall wealth.

Fuchs-Schündeln, 2008; Gertler et al., 2016).

Hence, we model the flow of savings in a particular period, denoted $\Delta Savings_{it}$ (where $Savings_{it}$ is beneficiary *i*'s stock of savings in period *t*), as a function of the stock of savings in the previous period and income in the current period. Adding individual and time-period fixed effects, we have

$$\Delta Savings_{it} = \lambda_i + \delta_t + \theta Savings_{i,t-1} + \gamma Income_{it} + \varepsilon_{it}.$$
(2)

Despite not have enough periods of post-treatment data for most beneficiaries to reach their buffer stocks, modeling the flow of savings in this way allows us to estimate the equilibrium buffer stock of savings. Models of precautionary saving predict that $\theta < 0$, since the amount of new savings decreases as the stock of savings approaches the target level. In order to identify the effects of the debit card on the savings rate over time, we interact the above terms with a dummy indicating treatment localities and time period dummies.

We are not actually able to implement the above model as specified because we are restricted to using bank account information. Instead, we estimate the change in net account balances as a function of lagged net balances and transfers deposited during the period. Under a set of testable assumptions, we can interpret the estimated coefficients on interactions with the treatment dummy as causal effects of the debit card on the flow of savings. Specifically, we need to assume that (i) there are no deposits into the account other than the transfer, (ii) the debit card does not affect other sources of income, and (iii) the debit card does not affect other non-account savings. The first two assumptions imply that the debit card can only affect savings out of transfers and not through other sources of income. The last assumption implies that any increase in savings in the bank account does not substitute for other forms of saving; an increase in bank savings constitutes an increase in total savings. Empirically we find that all three assumptions hold. First, as we have already shown in Figure IX, almost no beneficiaries deposit any funds in addition to the transfers into their savings accounts in any period. Second, using household survey panel data in Section 7, we find that the debit cards do not affect income. Third, using the household survey data, we find the same magnitude effect of the debit card on total savings as we do with administrative bank account data.

Incorporating all of the above changes to (2) and allowing the debit card's effect to vary over

time with the card, we obtain the following specification:

Net $Balance_{it} - Net \ Balance_{i,t-1} =$

$$\lambda_{i} + \sum_{k} \delta_{k} \mathbb{I}(k=t) + \sum_{k} \alpha_{k} T_{j(i)} \times \mathbb{I}(k=t) + \sum_{k} \theta_{k} Net \ Balance_{i,t-1} \times \mathbb{I}(k=t) + \sum_{k} \xi_{k} Net \ Balance_{i,t-1} \times \mathbb{I}(k=t) - \sum_{k} \gamma_{k} Transfers_{it} \times \mathbb{I}(k=t) + \sum_{k} \psi_{k} Transfers_{it} \times T_{j(i)} \times \mathbb{I}(k=t) + \varepsilon_{it}.$$
(3)

The main advantage of this specification over the reduced-form analysis presented in Section 4.1 is that it allows existing balances to influence the savings rate, enabling us to test the prediction from precautionary saving models that as a beneficiary accumulates savings and approaches her target buffer stock, her rate of saving decreases. This allows us to predict the equilibrium buffer stock, despite not having enough time periods after treatment for all beneficiaries to reach their buffer stocks. An additional advantage is that it controls for the amount of transfers in each period, which varies both across households and within households over time.¹⁸

We estimate the effect of the debit card on the savings rate from the above specification as

$$\hat{\Phi}_k \equiv (\hat{\alpha}_k + \hat{\xi}_k \omega_{k-1} + \hat{\psi}_k \mu_k) / \overline{Y}, \tag{4}$$

where ω_{k-1} is average lagged net balance and μ_k is average transfers in period k; \overline{Y} is average income. The numerator in (4) gives the difference between treatment and control in the flow of savings in pesos; the denominator divides by average income to obtain the savings rate.¹⁹

The right-hand side of the specification in (3) includes individual fixed effects and lagged net balance; in this case, since the dependent variable is a function of net balance, the assumption that

¹⁸Results are robust to excluding the $Transfer_{it}$ interaction terms; see Figure B.3. Because transfer amounts vary for a number of reasons, we control for them in the preferred specification. When there is an election, federal law requires Oportunidades to give the transfer in advance so that there is no payment close to the election month. In practice, this means that beneficiaries receive no payment in the bimester of the election and an additional payment in the preceding bimester. If a family does not comply with program conditions such as school attendance and health check-ups, the payment is suspended, but if the family returns to complying with the conditions, the missed payment is added into a future payment. Payments also vary systematically by time of year, as the program includes a school component that is not paid during the summer, and a school supplies component that is only paid during one bimester out of the year. Finally, changes in family structure affect the transfer amount because one child might age into or out of the program, for example.

¹⁹Average income is obtained from the 2009–10 wave of the ENCELURB household survey conducted by Oportunidades (described in Section 3). It is scaled to a four-month period to match the time period of the estimated effect of the debit card on the flow of savings.

the individual fixed effects are uncorrelated with the error does not hold, and the bias could be significant if the number of time periods is small (Nickell, 1981). To account for this, we use a system GMM estimator proposed by Blundell and Bond (1998) that is consistent for fixed T, large N and performs well in Monte Carlo simulations (Bun and Kiviet, 2006). We report in Figure Xa the GMM-estimated effects of the debit card on the savings rate from (4).²⁰ We also present simple estimates from a model without fixed effects, estimated by OLS, in Figure Xb.

The results in Figure X show that during the pre-treatment period, there is no difference between the treatment and control groups in the savings rate: $\hat{\Phi}_k = 0$ for all $k < \ell$. After receiving the card, there is an initial period of 9-12 months before the beneficiaries begin saving in the account: during this period, we continue to observe $\hat{\Phi}_k = 0$, i.e. no difference between the treatment and control groups in savings rates. In Section 5 we test different types of learning that could explain this delayed effect. After the learning period, however, we observe a substantial increase in the treatment group's savings rate relative to that of the control. The difference between the treatment and control groups in savings rates in the account is 5.4% of income after one year with the card. Models of precautionary savings predict that the savings rate should fall once a positive savings balance is achieved, with the savings rate dampened by a negative coefficient on lagged balance. Indeed, this is what we find: in the following period, the effect of the debit card on the savings rate is 3.7% of income, and in the last period before the control group also receives cards, it is 2.7%.²¹ Of course, if we had data for additional periods in time before our control group also received cards, we would expect to see the savings rate continue to fall toward zero as those with cards reach their precautionary savings targets; nevertheless, beneficiaries would have built up an economically meaningful buffer stock during this period.

Ideally, we would also like to measure the equilibrium buffer stock that beneficiaries accumulate. Since many beneficiaries are still accumulating savings after two years with the card, we do not have sufficient time periods to measure their equilibrium buffer stock. Instead, to predict the buffer stock they will accumulate, we note that once a beneficiary has reached her equilibrium buffer stock, Net $Balance_{it} = Net \ Balance_{i,t-1}$. Plugging this into (3) and solving for the equilibrium

²⁰The results are presented for wave 1, where—unlike wave 2—we have enough post-treatment periods for learning to take place and afterwards for the flow of savings to change as the stock of savings approaches the savings target.

²¹See Appendix D for a full discussion of the estimated coefficients. We show that the estimates align with the predictions from models of precautionary saving not only for $\hat{\Phi}_k$, as we've shown here, but also for the coefficients governing the dynamic dampening effect of the stock of savings, $\hat{\theta}_k$ and $\hat{\xi}_k$.

balance for those with a card, we obtain Net Balance = $(\delta + \alpha + (\gamma + \psi) Transfers)/(-\theta - \xi)$. Using averages for these coefficients from the post-learning periods, we predict that the average equilibrium buffer stock is 2470 pesos; to put this quantity in context, it equals 55% of beneficiaries' monthly income.²²

5 Learning

The time delay before a beneficiary begins saving after receiving her debit card suggests that learning might be occurring. We explore three specific kinds of learning: (1) *learning to trust* that the bank is a safe place to save; (2) *learning the technology*, i.e., learning how to save in banks or use ATMs, learning where ATMs are located, learning about transaction costs, etc.; and (3) *learning Oportunidades' eligibility rule* that beneficiaries who accumulate savings will not be dropped from the program.²³ In order to test these hypotheses, we complement the administrative Bansefi data with data from two beneficiary surveys: (1) the 2012 Payment Methods Survey and (2) the 2010 ENCASDU. Both are cross-sectional surveys of stratified random samples of urban Oportunidades beneficiaries; the ENCASDU oversampled in localities that received cards in early 2009. The surveys were designed by Oportunidades to learn more about the expansion of debit cards, and in both cases we restrict our analysis to the sample of respondents who received their benefits in savings accounts tied to debit cards at the time of the survey.²⁴

5.1 Learning to Trust

Beneficiaries might delay starting to save in order to build their trust that the bank is not reducing their account balances by charging hidden fees or through outright stealing. The debit card lowers the cost of checking account balances, leading to an increase in balance checks. Although a beneficiary could check her balance at Bansefi branches prior to receiving the card, the debit card makes it much more convenient since it allows balance checks at any bank's ATM.²⁵ We hypothesize that by checking her balance and seeing that the amount is as expected, the beneficiary learns that the bank

²²The 95% confidence interval of the equilibrium buffer stock is [2015, 2926], estimated using the delta method.

²³These potential learning explanations were conjectured in a meeting with Oportunidades program officials where we shared our savings results.

²⁴The questions we use were not asked to those who had not yet received debit cards.

²⁵The median household lives 5.2 kilometers (using the shortest road distance) from the nearest Bansefi branch, compared to 1.1 kilometers from an ATM.

is not stealing any money or applying hidden fees. In turn, the client updates downward her prior about the risk of losing money. With simple Bayesian learning, balance checking has decreasing marginal benefit as she updates her beliefs, which would lead to a decrease in the number of balance checks over time. Hence, over time with the card, we expect balance checks to fall and trust to rise. We test the hypothesis that balance checks fall over time with both the administrative and survey data. We then examine whether higher savings balances are correlated with the number of balance checks within accounts in the administrative account data and use the survey data to test whether self-reported trust in the bank increases over time with the card.

Balance Checks Over Time with the Debit Card

We first use the Bansefi transactions data to test whether, as we hypothesize, balance checks fall over time with the card.²⁶ Figure XI plots the average number of times clients check their balances over time in wave 1, with the dashed vertical line indicating the timing of card receipt. In the first period after receiving debit cards, beneficiaries check their balances 3 times on average, and the number of balance checks falls to 1.4 after two years with the card.

A possible explanation for this pattern is that clients are not checking if their existing balances remain, but if a new transfer has arrived. To reject this alternative explanation, we also plot in Figure XI the average number of balance checks, conditional on these checks occurring after the transfer was received in that same bimester, and on a different day than a withdrawal. Since this subset of balance checks happens after the transfer is received, and no money is withdrawn on that day (as would be expected if they were checking if the transfer had been deposited and discovered that it had), we posit that these checks occur precisely to monitor the bank and verify that existing balances remain in the account. We observe that immediately after receiving the card, recipients check their balances in this fashion 2.2 times on average per period, and only 0.4 times per period after two years with the card.²⁷

A final concern is that for learning to occur, clients must have a positive balance. Since we do

²⁶We do not observe balance checks at Bansefi branches in our transactions data since these are not charged a fee; hence, we do not observe balance checks prior to receiving the card. Nevertheless, it is unlikely that beneficiaries used this mechanism to monitor the bank prior to receiving a debit card due to the relatively high indirect costs of traveling to the nearest Bansefi branch.

²⁷Beneficiaries were given calendars with exact transfer dates and hence should know the dates on which transfers are deposited (see Figure A.3). We remain conservative by excluding all balance checks that *could be* interpreted as checking if a new transfer has arrived, i.e. excluding all balance checks that occurred prior to the transfer being deposited or after the transfer being deposited on the same day that money is withdrawn.

not have daily account balances, we take the conservative approach of defining a balance as positive if the cumulative transfer amounts minus the cumulative withdrawal amounts in the bimester is positive at the time of the balance check (this is a sufficient but not necessary condition for the balance to be positive). Among the subset of balance checks that occur after the transfer is deposited and not on the same day as a withdrawal, 89% of accounts have a positive balance at the time of checking. Focusing on the first period, when such balance checks are very frequent, we observe that many checks occur on small but positive balances: the 25th percentile of balances at the time of a balance check is 20 pesos, the median is 55 pesos, and the 75th percentile is 110 pesos. This supports the hypothesis that a beneficiary would initially leave a small balance in her account in order to be able to check her balance and confirm that it is as expected.

We cross-validate the above results from the administrative data by applying a similar test using the self-reported number of balance checks per period from the Payment Methods Survey of beneficiaries. We exploit variation in length of time with the debit card to test whether those who have had the card longer make less balance checks. Specifically, we split the sample by the median time with the card and estimate the following model:

$$y_i = \alpha + \gamma \mathbb{I}(\text{Card} \le \text{median time})_i + u_i,$$
(5)

where y_i is either (i) the self-reported number of balance checks over the past bimester; or (ii) the self-reported number of balance checks over the past bimester without withdrawing any money. Figure XIIa shows the results: both the number of balance checks and the number of balance checks without withdrawing decrease over time with the card. Those who have had the card for more than the median time (12 months) make 31% fewer trips to the ATM to check their balances without withdrawing money than those who have had the card for less time.

Correlation of Savings Balances and Balance Checks

We now test whether a negative correlation between balance checks and savings exists *within* accounts. We hypothesize that initially, when trust is low, beneficiaries do not yet save in the account, but do use the card to make a high number of balance checks. As they build trust, they make fewer

balance checks, and this is followed by an increase in the stock of savings. Specifically, we estimate

Net
$$Balance_{it} = \lambda_i + \sum_{c \neq 0} \eta_c \mathbb{I}(Checks_{it} = c) + \varepsilon_{it},$$
 (6)

where Net Balance_{it} is the net balance in account *i* at time *t*, the λ_i are account-level (i.e., beneficiary) fixed effects, and Checks_{it} is the number of balance checks in account *i* over period *t*, which we top code at 5 to avoid having many dummies for categories of high numbers of balance checks with few observations.²⁸ The η_c coefficients thus measure the within-account correlation between the stock of savings and number of balance checks, relative to the 0 balance checks (c = 0) category. Our hypothesis that beneficiaries increase their savings as they decrease their number of balance checks predicts that $\eta_c < 0$, and that η_c is decreasing (i.e., becoming more negative) in *c*.

Figure XIII shows the results from (6). As predicted, account balance is negatively correlated with number of balance checks. Although there is no difference (precisely estimated) in balances when beneficiaries make 0 vs. 1 balance check, all coefficients for the categories corresponding to more than one balance check are negative, large, and statistically significant. For example, in periods where beneficiaries make two balance checks, their savings average 180 pesos less than in periods when they make zero or one balance checks. Furthermore, for every additional balance check, net balances decrease, and this decrease is statistically significant. For example, in periods in which beneficiaries make 5 or more balance checks, their balances are on average 422 pesos lower than periods in which those same beneficiaries make zero or one check. This shows that as a beneficiary checks her balance less, she increases her savings balance.

Trust Over Time with the Debit Card

In this section we test the hypothesis that the longer beneficiaries have had the debit card, the higher their trust in the bank. We measure trust as follows. The ENCASDU survey asks: "Do you leave part of the monetary support from Oportunidades in your bank account?" If the response is no, the respondent is then asked: "Why don't you keep part of the monetary support from Oportunidades in your Bansefi savings account?" *Lack of trust* is captured by the answer "because if I do not take out all of the money I can lose what remains in the bank" (one of the pre-written responses to the

²⁸We do not include time fixed effects because the within-account changes in the stock of savings over time constitute precisely the variation we are exploiting. As always, ε_{it} are clustered at the bank branch level.

question) or a similar open-ended response related to not trusting the bank.²⁹ If the respondent provides a different reason for not saving in the account, or answers the first question "Yes" (i.e., saves in the account), we code *lack of trust* as 0. We then estimate (5) with lack of trust as the dependent variable, again exploiting the exogenous variation in the length of time beneficiaries have had the card. Because ENCASDU was conducted in the second half of 2010 before any beneficiary from wave 2 received cards, those with less and more time with the card in this comparison all received cards during wave $1.^{30}$

Figure XIV shows the results. Lack of trust is cited as the reason for not saving by 24% of those who have had the card for less than the median time. Trust increases over time, however, and beneficiaries with more than the median time with the card are 33% less likely to report not saving due to low trust.³¹

5.2 Learning the Technology

During the period of delay before starting to save, beneficiaries could be learning over time how to use their debit cards, learning that they can save in the account, learning where ATMs are located, or learning the transaction costs of using the account. To address the first two of these possibilities, the Payment Methods Survey includes various questions about use of the accounts after receiving debit cards: specifically, each respondent is asked whether (i) it is hard to use the ATM; (ii) she gets help using the ATM; and (iii) she knows her PIN. Thus, we estimate (5) with each of these three dependent dummy variables. Figure XIIb shows that there is no statistically significant difference between the group that has had the card for less than the median time compared to the group that has had the card more than the median time.

In the ENCASDU, we use the same direct survey question from Section 5.1 on self-reported reasons for not saving to test whether beneficiaries don't save due to lack of knowledge about how

²⁹The survey question allows the beneficiary to select one of the pre-written responses, or answer "other" and provide an open-ended response. 5% use the open-ended option. Examples of open-ended responses that were coded as lack of trust include "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."

 $^{^{30}}$ In Table B.1, we show balance between those who have had the card for more vs. less than the median time. We find no statistically significant differences at the 5% level and one statistically significant difference (out of 10 variables) at the 10% level, as would be expected by chance.

³¹Note that because of the timing of the ENCASDU survey, those with the card for less than the median time have nevertheless had the card for at least 9 months, meaning that some of them would have likely developed trust in the bank prior to being surveyed. Those with more than the median time with the card have had it for 5 months longer on average.

to save in the account. Lack of this type of knowledge, however, is rarely cited as a reason for not saving in the ENCADSU survey: less than 2% of beneficiaries cite not saving due to lack of knowledge, and there is no difference between who that have had the card for less than and more than the median time (Figure XIV).

In addition to finding little evidence of increased knowledge of the technology, we find that use of the accounts and ATMs increases immediately after receiving the card, then remains fairly stable over time. This is inconsistent with the hypothesis of learning where ATMs are located. Using the administrative data, we saw this pattern for withdrawals in Figure VIII; we can also test if clients immediately start using the card to withdraw at ATMs and convenience stores rather than bank branches. Figure XV shows the percentage of clients in each wave who use their debit card to make at least one withdrawal at an ATM or convenience store instead of going to the bank branch: depending on the bimester, 93–95% of clients use them to withdraw at ATMs and convenience stores and the adoption rate appears nearly instantaneous.³²

The *learning the technology* hypothesis is also inconsistent with the evolution of balance checks over time. As a beneficiary learns the technology, it should become easier (i.e., less costly) for her to check her balance. The fall in the marginal cost of using the ATM should then increase the number of balance checks over time. As shown in Section 5.1, however, we find the opposite trend in balance checks: the number of balance checks falls over time.

Finally, beneficiaries might be learning about ATM transaction costs, and start saving once they learn that these are sufficiently low. We test this alternative story directly using two questions from the Payment Methods Survey asking beneficiaries if they know how much the bank charges them for each (i) balance check and (ii) withdrawal after the initial free withdrawals. We find that the self-reported cost of transactions is not different for beneficiaries who have had the card for less vs. more than the median time: Figure XIIc displays, by time with the card, beneficiaries' self-reported estimates of these fees to check balances and withdraw. There is no difference in beneficiaries' self-reported estimates of transaction costs based on time with the card.³³

³²In the period in which cards are received, 82% of beneficiaries make a withdrawal at an ATM; this figure is lower than in subsequent periods, as expected, since some beneficiaries received the card mid-period after already receiving and withdrawing their benefits.

³³Beneficiaries are also fairly accurate. The median *actual* balance check fee in the transactions data is 10.4 pesos, while the median fee *estimated by beneficiaries* is 11 pesos; more importantly, these estimates do not vary by how long beneficiaries have had the card, as shown in Figure XII. The median withdrawal fee is 40 pesos, while the median estimated withdrawal fee is 24 pesos. While beneficiaries underestimate withdrawal fees (which are only charged after the second withdrawal in the bimester), the estimates do not differ by time with the card.

5.3 Learning the Program Rules

The last type of learning conjectured by Oportunidades program officials was that beneficiaries may have initially thought that saving in the account would make them be viewed as less poor and thus ineligible for the program, but learned over time that this was not the case. Due to this salient concern among program officials, the Payment Methods Survey includes the following pre-written response to the question about reasons for not saving: "because if I save in the account, they can drop me from Oportunidades."

We thus estimate (5) with the dependent variable equal to 1 if respondents do not save for this reason (which we call *fear of ineligibility* in Figure XIV), or a related reason listed in response to the optional open-ended response to the same survey question.³⁴ All other beneficiaries with savings accounts and debit cards are coded as 0 (including if they reported saving in the account in response to the previous survey question). The first thing to note from Figure XIV is that fear of being dropped from the program due to having savings in the bank is rarely cited as a reason for not saving, accounting for less than 4% of the sample who have had the card for less than the median amount of time. Furthermore, there is no statistically significant difference comparing these beneficiaries to those who have had the card for more than the median amount of time. This is consistent with information from our meetings with Oportunidades program officials, in which they reported that when initially providing bank accounts, they emphasized to beneficiaries that saving in the account would not disqualify them from future benefits.

6 The Relationship between Trust and Saving

In this section, we directly estimate the relationship between reported trust in the bank and the savings rate. In Section 5 we found that time with the card increases trust, but does not affect knowledge of how to use the technology, transaction costs, or program rules regarding saving in the account. Consistent with these results, we assume in this section that time with the card affects saving only through its effect on trust. If this assumption holds, we can express the reduced-form

³⁴Examples of open-ended responses coded as *fear of ineligibility* include "because they say that the card gets canceled if we don't withdraw the entire benefit" and "because they told me that if I don't take my benefit in a single withdrawal, the account would be frozen."

effect of time with the card on the flow of savings as

$$\frac{\mathrm{d}\Delta Savings}{\mathrm{d}Time \ with \ card} = \frac{\partial\Delta Savings}{\partial Trust} \cdot \frac{\partial Trust}{\partial Time \ with \ card.}$$
(7)

Our goal here is to directly estimate the relationship between reported trust in the bank and the savings rate, i.e., the first term of the right-hand side of (7). To do this, we merge the administrative data on net balances (from Section 4.4) with the ENCASDU survey data on trust (from Section 5.1). Everyone in this sample has had the card for between 9 and 18 months; we exploit variation in time with the card for identification. Since all of the beneficiaries in this sample have the card, all benefit from the lower transaction costs that debit cards engender. Using administrative identifiers provided by Oportunidades, we are able to merge 1330 of the 1694 beneficiaries in the survey with their corresponding administrative savings data; among this sample, we restrict the administrative savings data to the two bimesters that overlap with the timing of the survey.

To estimate the effect of trust on saving, we regress the flow of savings $\Delta Savings_{it}$ on a trust dummy (which is the complement of the *lack of trust* dummy used in Section 5.1):

Net
$$Balance_{it} - Net \ Balance_{i,t-1} = \zeta Trust_{it} + \varepsilon_{it}.$$
 (8)

In the OLS regression, we find no relationship between reported trust and the flow of savings. This is not surprising, as trust is endogenous: in the cross-section, those with initially high trust prior to the card or who developed trust in the bank quickly may have already reached their savings targets and thus not be adding additional savings. Furthermore, self-reported trust is likely measured with error. Since trust is endogenous and measured with error, we instrument it with the date of debit card assignment; this isolates the variation in trust that can be explained exogenously by time with the card. We already know from Section 5 that this instrument has a strong first stage.

Three pieces of evidence suggest that the instrument satisfies the exclusion restriction. First, time with the card is uncorrelated with sociodemographic characteristics. Second, time with the card does not affect other types of learning, as shown in Section 5. Third, *time* with the card (as opposed to the card itself) does not affect transaction costs, which are immediately reduced upon receiving the card: beneficiaries react instantaneously to the reduction in transaction costs by increasing the number of withdrawals (Figure VIII) and switching to withdrawing at ATMs (Figure XV). After

they make these immediate behavioral changes upon receiving the card, withdrawals per bimester and the proportion of withdrawals made at ATMs are constant over time. Recall that, in the sample used in this section, everyone has had a card for at least 9 months; transaction costs do not change as a result of having the card for additional months.

Table II reports the OLS and IV results from estimating (8), where in the IV regression trust is instrumented with a set of dummy variables for the timing of debit card receipt. Coefficients are expressed as a proportion of average income (from the survey) and standard errors are clustered at the locality level. The first stage, i.e. the effect of timing of debit card receipt on trust, has an F-statistic of 40. Taking a weighted average of the coefficients on each debit card timing dummy, the first stage shows that an average of six additional months with the card leads to a 10.3 percentage point increase in the probability of trusting the bank. The IV coefficient in column 2 shows that beneficiaries who report trusting the bank as a result of having the card for an additional six months save an additional 2.8% of their income, statistically significant at the 5% level.³⁵ The IV coefficient corresponds to the effect of being induced to trust the bank by virtue of having the debit card for a longer period of time (and hence having sufficient time to build trust in the bank).

To conclude, assuming that *additional* time with the card affects saving only through trust, we find a direct effect of trust on the flow of savings at the beneficiary level. A beneficiary who switches from not trusting the bank to trusting it as a result of having the card longer increases her savings rate by 2.8% of income.³⁶ This is, to our knowledge, the first direct causal estimate in the literature of the effect of trust in financial institutions on formal saving.

³⁵The results are robust to estimating a specification analogous to (3) based on models of precautionary saving, controlling for the lagged stock of savings and current transfers, interacted with trust. The instruments are again strong: the Sanderson and Windmeijer (2016) multivariate F-test for IV models with multiple endogenous variables (in this case, trust and its interactions) gives F-statistics of 18 for trust, 147 for trust interacted with lagged net balance, and 38 for trust interacted with transfers. The result in Table II column 3 shows that, using this alternative specification, the part of trust explained by the timing of debit card receipt accounts for a savings rate increase of 2.9% of income (significant at the 10% level), consistent with the results from the simpler specification.

 $^{^{36}}$ The direct effect of trust on saving is toward the lower end of the range of reduced-form effects we find in Section 4.4, which are from 2.7% to 5.4%. However, the evolution of average balances suggests a non-linear relationship between time with the card and trust, and the samples are not identical: while the results in Section 4.4 are based on a treatment group that has the card for up to two years versus a control group that has not yet received a card, in the ENCASDU sample everyone has had the card for at least 9 months. Indeed, estimating the reduced-form effect in this sample shows that an additional 6 months with the card (relative to having it for at least 9 months) increases savings by 0.3% of income.

7 Increase in Overall Savings vs. 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 saving. This section investigates whether the observed increase in Bansefi account savings crowds out other savings. We take advantage of Oportunidades' ENCELURB panel survey, conducted in urban and semi-urban localities in four waves during the years 2002, 2003, 2004 and November 2009 to 2010. This survey is conducted by Oportunidades and has comprehensive modules on consumption, income, and assets for 6272 households.³⁷ Of these, 2942 households live in urban areas and received their benefits in a savings account prior to the rollout of debit cards; these households make up our sample.

We use a simple difference-in-differences identification strategy where we examine changes in consumption, income, saving, purchases of durables, and the stock of assets across beneficiaries, exploiting the differential timing of debit card receipt. Because the ENCELURB was conducted after wave 1 localities had received cards but before wave 2 or control localities had received cards, we compare those with cards (wave 1) to those who had not yet received cards (wave 2 and control), respectively referring to them as "treatment" and "control" localities in this section. The identification assumption is that in the absence of the debit card, treatment and control groups would have experienced similar changes in consumption, income, saving, and assets. We formally test for parallel pre-treatment trends for each of our dependent variables in Table B.4 and fail to reject the null hypothesis of parallel trends.

Having established that the identification assumption is plausible, we estimate

$$y_{it} = \lambda_i + \delta_t + \gamma D_{j(i)t} + \nu_{it}, \tag{9}$$

separately for five dependent variables: consumption, income, flow of savings (constructed as income

³⁷The 2002, 2003, and 2004 waves had around 17,000 households, but due to budget constraints the number of localities was cut for the 2009–10 wave. We restrict our sample to those included in the 2009–10 wave; not every household was surveyed in every baseline wave, resulting in an unbalanced sample. The consumption, income, and assets modules of Oportunidades' analogous survey for rural areas have been used by Angelucci and De Giorgi (2009), Gertler, Martinez, and Rubio-Codina (2012), Attanasio et al. (2013), and de Janvry et al. (2015), while these modules from the ENCELURB have been used by Behrman et al. (2012) and Angelucci and Attanasio (2013), among others.

minus consumption), purchase of durables, and an asset index.³⁸ All variables except the asset index are measured in pesos per month, *i* indexes households, and *t* indexes survey rounds.³⁹ Variables are winsorized at the 5% level to avoid results driven by outliers. Time-invariant differences in household observables and unobservables are captured by the household fixed effect 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 household *i* lived prior to treatment has received debit cards by time *t*. We use the locality of residence prior to treatment to avoid confounding migration effects, and estimate cluster-robust standard errors clustered by locality.

If the increase in formal savings is merely a substitution away from other forms of saving, we expect to find $\gamma = 0$ when the dependent variable is the flow of total savings (defined as income minus consumption). And if the form of savings that beneficiaries substituted away from was durable assets, we expect $\gamma < 0$ for the stock of assets, and potentially also for the purchase of durables. If, on the other hand, the formal savings increase constitutes an increase in total savings, then we expect $\gamma > 0$ for the flow of total savings; if there is partial crowding out, we expect the magnitude of γ to be less than the magnitude found in the administrative Bansefi data, while if there is no crowding out, we expect the magnitude to be similar. Furthermore, one of the assumptions in Section 4.4 was that the debit card does not affect income, so we test $\gamma = 0$ for income. After confirming there is no effect on income, we expect $\gamma < 0$ for consumption, since consumption must decrease if total savings increases and income does not change. Furthermore, if there is no substitution of savings from assets (and if they are not using the formal savings accounts to save up for assets, at least in the short run), we expect $\gamma = 0$ for the purchase of durables (which measures a flow) and the asset index (which measures a stock).

Our findings indicate that the increase in formal savings shown in Section 4 represents an increase in overall savings. Figure XVI shows that consumption decreased by about 138 pesos per month on average (statistically significant at the 5% level). We do not find any effect on income.⁴⁰ Purchases of durables and the stock of assets do not change, ruling out a crowding out of these forms of saving.

³⁸Our measure of the flow of savings is imperfect, but is commonly used in the literature (e.g., Dynan, Skinner, and Zeldes, 2004).

³⁹The asset index dependent variable is constructed as the first principal component of dummy variables indicating ownership of the assets that are included in all rounds of the survey questionnaire: car, truck, motorcycle, TV, video or DVD player, radio, washer, gas stove and refrigerator.

⁴⁰We also test the difference in the coefficients of consumption and income using a stacked regression (which is equivalent to seemingly unrelated regression when the same regressors are used in each equation, as is the case here); although both are noisily measured, the difference in the coefficients is significant at the 10% level (p = 0.092).

The increase in the flow of savings, measured as income minus consumption, is estimated at 236 pesos per month, and is statistically significant at the 5% level. These results are robust to the extent of winsorizing and to allowing flexible time trends as a function of household characteristics.⁴¹

These results mean that total savings—not just account savings—increase, and that this increase in being funded by lower consumption today. A back-of-the-envelope calculation reveals that the magnitude of the increase in the flow of savings from the household survey data is about the same as that of the increase in the flow of savings in the Bansefi account. In Section 4.4 we estimate that after 1 year with the card, beneficiaries who received cards in wave 1 save 5.4% more of their income than the control group. In our survey data, the effect we find on the flow of total savings is 236 pesos per month; dividing by average household income in the post-treatment survey wave, 4,629 pesos per month, equates to 5.1% of income. We cannot reject that the effect sizes in the administrative data and survey data are equal; since the point estimate from the survey data has a large confidence interval, however, the more convincing result is that the point estimate of the savings effect from the survey data is within the tight [0.045, 0.062] confidence interval of the effect size from the more precise administrative data. These results suggest that most or all of the increase in savings in the account is new savings, and that there is no crowd-out of other types of saving.

These results are consistent with Dupas and Robinson (2013a), Ashraf, Karlan, and Yin (2015), and Kast, Meier, and Pomeranz (2016), who find that offering formal savings accounts does not crowd out other forms of saving. This may be because saving informally is difficult for a number of reasons, so access to a trusted formal savings account allows households to achieve a higher level of overall savings. First, informal savings can be stolen (Banerjee and Duflo, 2007; Schechter, 2007; Alvarez and Lippi, 2009). Second, intra-household bargaining issues may prevent women from saving at home (Anderson and Baland, 2002; Ashraf, 2009; Schaner, 2015). Third, money saved at home could be in demand from friends and relatives (Baland, Guirkinger, and Mali, 2011; Dupas

⁴¹Table B.2 shows that the effects are robust to using the raw data without winsorizing (column 1) and to winsorizing at 1% (column 2) or 5% (column 3, which are our main results presented in Figure XVI); we follow Kast and Pomeranz (2014) who show the robustness of results to these three possibilities for their savings measures. They are also robust to including baseline characteristics interacted with time fixed effects (column 4). The baseline characteristics that we interact with time fixed effects in column 4 include characteristics of the household head (working status, a quadratic polynomial in years of schooling, and a quadratic polynomial in age), whether anyone in the household has a bank account, a number of characteristics used by the Mexican government to target social programs (the proportion of household members with access to health insurance, the proportion age 15 and older that are illiterate, the proportion ages 15-29 with less than 9 years of schooling), and dwelling characteristics (dirt floors, no bathroom, no piped water, no sewage, and number of occupants per room).

and Robinson, 2013b; Jakiela and Ozier, 2016). Finally, it may be tempting to spend money that had been intended to be saved if it is easily accessible, especially if the beneficiary has easy access to money saved informally at times when she is more financially constrained (Carvalho, Meier, and Wang, 2016). Once the bank is trusted, the account might form a soft commitment device that overcomes these self-control problems (Ashraf, Karlan, and Yin, 2006b; Bryan, Karlan, and Nelson, 2010).

8 Alternative Explanations

We have argued that the card allows beneficiaries to build trust in the bank by monitoring the bank's activity through balance checks. We now explore alternative explanations for the observed delayed effect, followed by a gradual increase in the savings balance and a change in the savings rate that adheres to predictions from models of precautionary saving and saving to purchase durables.

8.1 Supply-Side Expansion

An alternative explanation for the delayed effect and increase in savings over time is that banks gradually expanded complementary infrastructure (e.g., the number of ATMs) in localities where treated beneficiaries live, potentially as an optimal response to an increase in demand for financial services from the new cardholders in those localities. More ATMs would decrease the transaction cost of accessing funds, which could boost savings once the transaction cost is low enough that the bank becomes a desirable place to save (since, in precautionary savings models, a shock forces the client to incur the cost of an additional trip to the ATM). This explanation, which would imply a delayed decrease in transaction costs for some beneficiaries, is inconsistent with the immediate increase in the number of withdrawals we observe in Figure VIII.

We nevertheless directly test this hypothesis using quarterly data on the number of ATMs at the municipality level to see if there was a contemporaneous expansion of infrastructure that was correlated geographically with Oportunidades debit card expansion. Specifically, we estimate

$$y_{mt} = \lambda_m + \delta_t + \sum_{k=-6}^{6} \beta_k D_{m,t+k} + \varepsilon_{jt}, \qquad (10)$$

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} equals one if at least one locality in municipality m has Oportunidades debit cards in quarter t. The error term ε_{jt} is clustered by municipality. We include one and a half years (six quarters) of lags to test whether the supply of ATMs or bank branches responds to the rollout of debit cards, which from the perspective of banks can be thought of as a discrete jump in the number of potential users. We also include six quarters of leads to test whether the rollout of debit cards instead followed an expansion of bank infrastructure, which would be a threat to validity.

We use data on the number of ATMs and bank branches by bank by municipality by quarter from the Comisión Nacional Bancaria y de Valores (CNBV), from the last quarter of 2008—the first quarter for which data are available—through the last quarter of 2011, which is the end of our study period. We separately test whether lags of debit card receipt predict banking infrastructure (i.e., whether there is a supply-side response to the rollout of debit cards) by testing $\beta_{-6} = \cdots = \beta_{-1} = 0$, and 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) by testing $\beta_1 = \cdots = \beta_6 = 0$. We find evidence of neither relationship, failing to reject the null hypotheses of zero correlation between the rollout of debit cards and the expansion of banking infrastructure for each of the four dependent variables (Table B.3).⁴²

8.2 Local Income Shocks

A second alternative explanation is that the increase in savings is due to local macro shocks to incomes at the locality level. Given the geographical breadth of the treatment and control groups throughout Mexico, however, this is unlikely. Furthermore, if this were the case we would expect to find a differential change in income between the treatment and control groups after treatment; we directly test this hypothesis in Section 7 and find no differential change in income after treatment.

8.3 Time with the Bank Account

Finally, we investigate whether individuals learn about banks in general the longer they have their savings account itself (regardless of whether they have a debit card). This could only explain

⁴²This lack of a supply-side response by private banks is not illogical: the banks would have to make sufficient profit off of the new cardholders to justify the cost of installing new ATMs. Oportunidades beneficiaries with debit cards only make 1.4 withdrawals per bimester on average, and may not constitute a large enough share of the population in urban localities to justify the cost of installing new ATMs.

the differential savings trends in treatment and control groups if the debit card were a necessary condition for saving, but learning about the bank through having the *account* (not card) for a sufficient amount of time were also a necessary condition.

There are a number of reasons why experience with the savings account rather than time with the debit card itself cannot explain the savings effect. First, because the savings accounts were rolled out between 2002 and 2004 (Figure III), beneficiaries had already experienced several years with the account by 2009, when debit cards were first introduced. Indeed, the median month of account opening is October 2004, and less than 5% of accounts had existed for less than two years before they received debit cards. Second, both treatment and control accounts are accumulating time with their savings accounts simultaneously, and they have had accounts for the same amount of time on average. Third, our results from Section 4.4 include account fixed effects, so any time-invariant effect of having the account for a longer period of time would be absorbed. Fourth, we test whether results on savings rates vary when we split the sample based on whether the account was opened before or after the median date in Figure B.4. We find similar results across the two subsamples.

9 Conclusion

Trust in financial institutions is low, especially among the poor, and this may be a barrier to financial inclusion. A lack of trust in banks could explain why a number of studies offering the poor savings accounts with no fees or minimum balance requirements have found low take-up and, even among adopters, low use of the accounts (e.g., Dupas et al., 2016b). We show that the trust barrier is not insurmountable: it can be overcome by debit cards, a scalable existing technology. Our first important result is that debit cards have a large causal effect on savings in the account. Our second result is that increasing trust plays an important role in the savings effect. Once beneficiaries build trust in banks by using their debit cards to repeatedly check account balances, they begin to save and their savings increase over time. Our third result is that savings in the bank account are new savings, rather than a substitution from other forms of saving.

It is worth noting that beneficiaries with the debit card voluntarily use the technology and build savings in the account (whereas they could continue withdrawing all of their benefits from the bank branch, as they did prior to receiving the card), indicating a revealed preference for saving in formal financial institutions after building trust. Because the formal accounts pay no interest, this action also reveals an unmet demand for savings products among program beneficiaries.

The size of the savings effect—between 3 and 5% of income after 1–2 years with the card—is larger than that of various other savings interventions, including offering commitment devices, nofee accounts, higher interest rates, lower transaction costs, and financial education. Figure XVII compares the size of the effect we find on the savings rate to the effect of various other savings interventions studied in the literature.⁴³ One other intervention found to have a similarly large effect is mobile money, which is another technology that enables clients to easily check account balances and build trust. Suri and Jack (2016) find that female-headed households increase their savings rate by 3% of income on average, after six years of mobile money exposure. The large savings rate effect we estimate can be explained by the fact that once beneficiaries trust the bank, they start building precautionary savings. Extrapolating our estimates from the precautionary savings model to future periods, we predict that beneficiaries are saving towards an equilibrium buffer stock of 2470 pesos on average, which corresponds to 55% of their monthly income.

These results are important for public policy, as building savings in formal financial institutions has been shown to have positive welfare effects for the poor by enabling them to decrease consumption volatility (Chamon, Liu, and Prasad, 2013; Prina, 2015), accumulate money for microenterprise investments (Dupas and Robinson, 2013a), invest in preventative health products and pay for unexpected health emergencies (Dupas and Robinson, 2013b), invest in children's education (Prina, 2015), increase future agricultural/business output and household consumption (Brune et al., 2016), and decrease debt (Atkinson et al., 2013; Kast, Meier, and Pomeranz, 2016). For these reasons, Mullainathan and Shafir (2009) conclude that access to formal savings services "may provide an important pathway out of poverty."

Interventions that enable account holders to monitor banks and increase their trust in financial institutions may be a promising avenue to enable the poor to save in the formal financial sector. These interventions take advantage of prevalent technologies—such as debit cards, ATMs, point of sale terminals, and mobile phones. Governments and non-governmental organizations are increasingly using these technologies to digitize their social cash transfer programs, providing the opportunity to rapidly scale these trust-building technologies and enable the poor to save more.

⁴³We restrict this comparison to studies with a duration of at least 6 months. The median study duration is 14 months. See Appendix E for details on each study and how we use the results from the study's tables or replication data to estimate the impact of the savings intervention on the savings rate.

PRINCETON UNIVERSITY UNIVERSITY OF CALIFORNIA, BERKELEY, and NBER UNIVERSITY OF CALIFORNIA, BERKELEY ITAM

References

- Aker, Jenny C., Rachid Boumnijel, Amanda McClelland, and Niall Tierney, "Payment mechanisms and antipoverty programs: Evidence from a mobile money cash transfer experiment in Niger," Economic Development and Cultural Change, 65 (2016), 1–37.
- Algan, Yann and Pierre Cahuc, "Inherited trust and growth," American Economic Review, 100 (2010), 2060–2092.
- Alvarez, Fernando and Francesco Lippi, "Financial innovation and the transactions demand for cash," Econometrica, 77 (2009), 363–402.
- Anderson, Siwan and Jean-Marie Baland, "The economics of roscas and intrahousehold resource allocation," Quarterly Journal of Economics, 117 (2002), 963–995.
- Angelucci, Manuela and Orazio Attanasio, "The demand for food of poor urban Mexican households: Understanding policy impacts using structural models," American Economic Journal: Economic Policy, 5 (2013), 146–178.
- Angelucci, Manuela and Giacomo De Giorgi, "Indirect effects of an aid program: How do cash transfers affect ineligibles' consumption?" American Economic Review, 99 (2009), 486–508.
- Arellano, Manuel and Olympia Bover, "Another look at the instrumental variable estimation of error-component models," Journal of Econometrics, 68 (1995), 29–51.
- Arrow, Kenneth, "Gifts and exchanges," Philosophy & Public Affairs, 1 (1972), 343–362.
- Ashraf, Nava, "Spousal control and intra-household decision making: An experimental study in the Philippines," American Economic Review, 99 (2009), 1245–1277.
- Ashraf, Nava, Dean Karlan, and Wesley Yin, "Deposit collectors," Advances in Economic Analysis & Policy, 6 (2006a), 635–672.
- Ashraf, Nava, Dean Karlan, and Wesley Yin, "Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines," Quarterly Journal of Economics, 121 (2006b), 635-672.
- Ashraf, Nava, Dean Karlan, and Wesley Yin, "Savings in transnational households: A field experiment among migrants from El Salvador," Review of Economics and Statistics, 97 (2015), 332–351.
- Atkinson, Jesse, Alain de Janvry, Craig McIntosh, and Elisabeth Sadoulet, "Prompting microfinance borrowers to save: A field experiment from Guatemala," Economic Development and Cultural Change, 62 (2013), 21–64.
- Attanasio, Orazio, Vincenzo Di Maro, Valérie Lechene, and David Phillips, "Welfare consequences of food prices increases: Evidence from rural Mexico," Journal of Development Economics, 104 (2013), 136–151.
- Attanasio, Orazio, Luca Pellerano, and Sandra Polania Reyes, "Building trust? Conditional cash transfer programmes and social capital," Fiscal Studies, 30 (2009), 139–177.
- Attanasio, Orazio P., Luigi Guiso, and Tullio Jappelli, "The demand for money, financial innovation, and the welfare cost of inflation: An analysis with household data," Journal of Political Economy, 110 (2002), 317–351.
- Baland, Jean-Marie, Catherine Guirkinger, and Charlotte Mali, "Pretending to be poor: Borrowing to escape forced solidarity in Cameroon," Economic Development and Cultural Change, 60 (2011), 1–16.

- Banerjee, Abhijit V. and Esther Duflo, "The economic lives of the poor," Journal of Economic Perspectives, 21 (2007), 141–167.
- Baumol, William J., "The transactions demand for cash: An inventory theoretic approach," Quarterly Journal of Economics, 66 (1952), 545–556.
- Beaman, Lori, Dean Karlan, and Bram Thuysbaert, "Saving for a (not so) rainy day: A randomized evaluation of savings groups in Mali," NBER Working Paper 21169, 2014.
- Beck, Thorsten and Ross Levine, "Stock markets, banks, and growth: Panel evidence," Journal of Banking & Finance, 28 (2004), 423-442.
- Behrman, Jere R., Jorge Gallardo-García, Susan W. Parker, Petra E. Todd, and Viviana Vélez-Grajales, "Are conditional cash transfers effective in urban areas? Evidence from Mexico," Education Economics, 20 (2012), 233–259.
- Blundell, Richard and Stephen Bond, "Initial conditions and moment restrictions in dynamic panel data models," Journal of Econometrics, 87 (1998), 115–143.
- Bohnet, Iris, Benedikt Herrmann, and Richard Zeckhauser, "Trust and the reference points for trustworthiness in Gulf and Western countries," Quarterly Journal of Economics, 125 (2010), 811–828.
- Bold, Chris, David Porteous, and Sarah Rotman, "Social cash transfers and financial inclusion: Evidence from four countries," Consultative Group to Assist the Poor (CGAP) Focus Note 77, 2012.
- Breza, Emily and Arun G. Chandrasekhar, "Social networks, reputation and commitment: Evidence from a savings monitors experiment," NBER Working Paper 21169, 2015.
- Brune, Lasse, Xavier Giné, Jessica Goldberg, and Dean Yang, "Facilitating savings for agriculture: Field experimental evidence from Malawi," Economic Development and Cultural Change, 64 (2016), 187–220.
- Bryan, Gharad, Dean Karlan, and Scott Nelson, "Commitment devices," Annual Review of Economics, 2 (2010), 671–698.
- Bun, Maurice J. G. and Jan F. Kiviet, "The effects of dynamic feedbacks on LS and MM estimator accuracy in panel data models," Journal of Econometrics, 132 (2006), 409–444.
- Callen, Michael, Suresh de Mel, Craig McIntosh, and Christopher Woodruff, "What are the headwaters of formal savings? experimental evidence from Sri Lanka," NBER Working Paper 20736, 2014.
- Carroll, Christopher D., "Buffer-stock saving and the life cycle/permanent income hypothesis," Quarterly Journal of Economics, 112 (1997), 1–55.
- Carvalho, Leandro S., Stephan Meier, and Stephanie W. Wang, "Poverty and economic decisionmaking: Evidence from changes in financial resources at payday," American Economic Review, 106 (2016), 260–284.
- Chamon, Marcos, Kai Liu, and Eswar Prasad, "Income uncertainty and household savings in China," Journal of Development Economics, 105 (2013), 164–177.
- Chetty, Raj, John N. Friedman, SLeth-Petersen, Torben Heien Nielsen, and Tore Olsen, "Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark," Quarterly Journal of Economics, 129 (2014), 1141–1219.
- CNBV, "Comisión Nacional Bancaria y de Valores comunicado de prensa 094/2014,", 2014, URL http://www.cnbv.gob.mx/SECTORES-SUPERVISADOS/SECTOR-POPULAR/CasoFicrea/ Comunicado_de_Prensa_Revocaci%C3%B3n_FICREA.pdf.
- Cohn, Alain, Ernst Fehr, and Michel André Maréchal, "Business culture and dishonesty in the banking industry," Nature, 516 (2014), 86–89.
- de Janvry, Alain, Kyle Emerick, Marco Gonzalez-Navarro, and Elisabeth Sadoulet, "Delinking land rights from land use: Certification and migration in Mexico," American Economic Review, 105
(2015), 3125 - 3149.

- de Mel, Suresh, Craig McIntosh, and Christopher Woodruff, "Deposit collecting: Unbundling the role of frequency, salience, and habit formation in generating savings," American Economic Review: Papers & Proceedings, 103 (2013), 387–392.
- Deaton, Angus, "Saving and liquidity constraints," Econometrica, 59 (1991), 1221–1248.
- Deaton, Angus, "Household saving in LDCs: Credit markets, insurance and welfare," Scandinavian Journal of Economics, 94 (1992), 253–273.
- Drexler, Alejandro, Greg Fischer, and Antoinette Schoar, "Keeping it simple: Financial literacy and rules of thumb," American Economic Journal: Applied Economics, 6 (2014), 1–31.
- Dupas, Pascaline, Sarah Green, Anthony Keats, and Jonathan Robinson, "Challenges in banking the rural poor: Evidence from Kenya's Western Province," in Modernization and Development, Sebastian Edwards, Simon Johnson, and David N. Weil, eds. (Chicago: University of Chicago Press, 2016a).
- Dupas, Pascaline, Dean Karlan, Jonathan Robinson, and Diego Ubfal, "Banking the unbanked? Evidence from three countries," NBER Working Paper 22463, 2016b.
- Dupas, Pascaline and Jonathan Robinson, "Savings constraints and microenterprise development: Evidence from a field experiment in Kenya," American Economic Journal: Applied Economics, 5 (2013a), 163–192.
- Dupas, Pascaline and Jonathan Robinson, "Why don't the poor save more? Evidence from health savings experiments," American Economic Review, 103 (2013b), 1138–1171.
- Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes, "Do the rich save more?" Journal of Political Economy, 112 (2004), 397–444.
- Fuchs-Schündeln, Nicola, "The response of household saving to the large shock of German reunification," American Economic Review, 98 (2008), 1798–1828.
- Gertler, Paul, Orie Shelef, Catherine D. Wolfram, and Alan Fuchs, "The demand for energy-using assets among the world's rising middle classes," American Economic Review, 106 (2016), 1366–1401.
- Gertler, Paul J., Sebastian W. Martinez, and Marta Rubio-Codina, "Investing cash transfers to raise long-term living standards," American Economic Journal: Applied Economics, 4 (2012), 164–192.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales, "The role of social capital in financial development," American Economic Review, 49 (2004), 526–556.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales, "Trusting the stock market," Journal of Finance, 63 (2008), 2557–2600.
- Haushofer, Johannes and Jeremy Shapiro, "The short-term impact of unconditional cash transfers to the poor: Experimental evidence from Kenya," Quarterly Journal of Economics, 131 (2016), 1973–2042.
- Hirschman, Albert O., "Against parsimony: Three easy ways of complicating some categories of economic discourse," American Economic Review: Papers & Proceedings, 74 (1984), 89–96.
- Iyer, Rajkamal and Manju Puri, "Understanding bank runs: The importance of depositor-bank relationships and networks," American Economic Review, 102 (2012), 1414–1445.
- Jack, William, Adam Ray, and Tavneet Suri, "Transaction networks: Evidence from mobile money in Kenya," American Economic Review: Papers & Proceedings, 103 (2013), 356–361.
- Jack, William and Tavneet Suri, "Risk sharing and transactions costs: Evidence from Kenya's mobile money revolution," American Economic Review, 104 (2014), 183–223.
- Jakiela, Pamela and Owen Ozier, "Does Africa need a rotten kin theorem? Experimental evidence from village economies," Review of Economic Studies, 83 (2016), 231–268.
- Karlan, Dean, Margaret McConnell, Sendhil Mullainathan, and Jonathan Zinman, "Getting to the top of mind: How reminders increase saving," Management Science, 62 (2016), 3393–3411.

- Karlan, Dean, Markus Mobius, Tanya Rosenblat, and Adam Szeidl, "Trust and social collateral," Quarterly Journal of Economics, 124 (2009), 1307–1361.
- Karlan, Dean, Aishwarya Lakshmi Ratan, and Jonathan Zinman, "Savings by and for the poor: A research review and agenda," Review of Income and Wealth, 60 (2014), 36–78.
- Karlan, Dean and Jonathan Zinman, "Price and control elasticities of demand for savings," Working Paper, 2016, URL http://karlan.yale.edu/sites/default/files/savingselasticities_ dec2016_0.pdf.
- Kast, Felipe, Stephan Meier, and Dina Pomeranz, "Saving more in groups: Field experimental evidence from Chile," Harvard Business School Working Paper 12-060, 2016, URL http://www.hbs. edu/faculty/Publication%20Files/12-060_8c16f5e7-6fa1-48cc-858d-bba5f12c28ba.pdf.
- Kast, Felipe and Dina Pomeranz, "Saving more to borrow less: Experimental evidence from access to formal savings accounts in Chile," Harvard Business School Working Paper 14-001, 2014, URL http://www.hbs.edu/faculty/Publication%20Files/14-001_ 2237a8d2-6147-4667-9396-cd5ae9702f03.pdf.
- Knack, Stephen and Philip Keefer, "Does social capital have an economic payoff? A cross-country investigation," Quarterly Journal of Economics, 112 (1997), 1251–1288.
- La Porta, Rafael, Florencio López-de-Silanes, Andrei Shleifer, and Robert W. Vishny, "Trust in large organizations," American Economic Review: Papers & Proceedings, 87 (1997), 333–338.
- La Porta, Rafael, Florencio López-de-Silanes, and Guillermo Zamarripa, "Related lending," Quarterly Journal of Economics, 118 (2003), 231–268.
- Levy, S. and N. Schady, "Latin America's social policy challenge: Education, social insurance, redistribution," Journal of Economic Perspectives, 27 (2013), 193–218.
- McMillan, John and Christopher Woodruff, "Interfirm relationships and informal credit in Vietnam," Quarterly Journal of Economics, 114 (1999), 1285–1320.
- Modigliani, Franco, "Life cycle, individual thrift, and the wealth of nations," American Economic Review, 76 (1986), 297–313.
- Mullainathan, Sendhil and Eldar Shafir, "Savings policy and decision-making in low-income households," in Insufficient Funds: Savings, Assets, Credit and Banking Among Low-Income Households, 121–145 (New York City: Russell Sage Foundation Press, 2009).
- Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar, "Building state capacity: Evidence from biometric smartcards in India," American Economic Review, 106 (2016), 2895–2929.
- Nickell, Stephen, "Biases in dynamic models with fixed effects," Econometrica, 49 (1981), 1417–1426.
- Parker, Susan W. and Petra E. Todd, "Conditional cash transfers: The case of Progresa/Oportunidades," Journal of Economic Literature, (forthcoming).
- Prina, Silvia, "Banking the poor via savings accounts: Evidence from a field experiment," Journal of Development Economics, 115 (2015), 16–31.
- Roodman, David, "A note on the theme of too many instruments," Oxford Bulletin on Economics and Statistics, 71 (2009), 135–158.
- Rosenzweig, Mark R., "Savings behaviour in low-income countries," Oxford Review of Economic Policy, 17 (2001), 40–54.
- Rosenzweig, Mark R. and Kenneth I. Wolpin, "Credit market constraints, consumption smoothing, and the accumulation of durable production assets in low-income countries: Investments in bullocks in India," Journal of Political Economy, 101 (1993), 223–244.
- Sanderson, Eleanor and Frank Windmeijer, "A weak instrument F-test in linear IV models with multiple endogenous variables," Journal of Econometrics, 190 (2016), 212 – 221.
- Sapienza, Paola and Luigi Zingales, "A trust crisis," International Review of Finance, 12 (2012), 123–131.
- Sayinzoga, Aussi, Erwin H. Bulte, and Robert Lensink, "Financial literacy and financial behaviour:

Experimental evidence from rural Rwanda," Economic Journal, 126 (2016), 1571–1599.

- Schaner, Simone, "Do opposites detract? Intrahousehold preference heterogeneity and inefficient strategic savings," American Economic Journal: Applied Economics, 7 (2015), 135–174.
- Schaner, Simone, "The persistent power of behavioral change: Long-run impacts of temporary savings subsidies for the poor," NBER Working Paper 22534, 2016.
- Schaner, Simone, "The cost of convenience? Transaction costs, bargaining, and savings account use in Kenya," Journal of Human Resources, (forthcoming).
- Schechter, Laura, "Theft, gift-giving, and trustworthiness: Honesty is its own reward in rural Paraguay," American Economic Review, 97 (2007), 1560–1582.
- Seshan, Ganesh and Dean Yang, "Motivating migrants: A field experiment on financial decisionmaking in transnational households," Journal of Development Economics, 108 (2014), 119–127.
- Skinner, Jonathan, "Risky income, life cycle consumption, and precautionary savings," Journal of Monetary Economics, 22 (1988), 237–255.
- Somville, Vincent and Lore Vandewalle, "Saving by default: Evidence from a field experiment in rural india," Working Paper, 2016.
- Suri, Tavneet and William Jack, "The long-run poverty and gender impacts of mobile money," Science, 354 (2016), 1288–1292.
- Suri, Tavneet, William Jack, and Thomas M. Stoker, "Documenting the birth of a financial economy," Proceedings of the National Academy of Sciences, 109 (2012), 10257–10262.
- Tobin, James, "The interest-elasticity of transactions demand for cash," Review of Economics and Statistics, 38 (1956), 241–247.
- Williamson, Oliver E., "Calculativeness, trust, and economic organization," Journal of Law and Economics, 36 (1993), 453-486.
- Windmeijer, Frank, "A finite sample correction for the variance of linear efficient two-step GMM estimators," Journal of Econometrics, 126 (2005), 25–51.
- Ziliak, James P., "Efficient estimation with panel data when instruments are predetermined: An empirical comparison of moment-condition estimators," Journal of Business & Economic Statistics, 15 (1997), 419-431.
- Zinman, Jonathan, "Credit or debit?" Journal of Banking and Finance, 33 (2009), 358-366.



Figure I: Low Trust in Banks by Education Level in Mexico

Source: World Values Survey, Mexico, Wave 6 (2012).

Notes: N = 1993 individuals. Low trust in banks is defined as "not very much confidence" or "none at all" for the item "banks" in response to the following question: "I am going to name a number of organizations. For each one, could you tell me how much confidence you have in them: is it a great deal of confidence, quite a lot of confidence, not very much confidence or none at all?" Whiskers denote 95% confidence intervals.



Figure II: Cross-Country Comparison of Trust in Banks and Saving in Financial Institutions

Sources: World Values Survey (WVS), Wave 6 (2010–2014); Global Findex; World Development Indicators (WDI). Notes: N = 56 countries. The y-axis plots residuals from a regression of the proportion that save in financial institutions (from Global Findex) against controls (average age, education, and perceived income decile from WVS, GDP per capita and growth of GDP per capita from WDI). The x-axis plots residuals from a regression against the same controls of the proportion that respond "a great deal of confidence" or "quite a lot of confidence" in response to the WVS question "could you tell me how much confidence you have in banks: a great deal of confidence, quite a lot of confidence, not very much confidence or none at all?" The solid line shows a kernel-weighted local polynomial regression, while the gray area shows its 95% confidence interval.



(a) Administrative Bank Account Data



Source: Number of Oportunidades bank accounts with cards and without cards by bimester is from administrative data provided by Oportunidades.

Figure IV: Evolution of Average Balances



Sources: Administrative data from Bansefi on average account balances by bimester and timing of card receipt. Notes: N = 7,895,620 account-bimester observations from 343,204 accounts over 24 bimesters. Average balances are winsorized at the 95th percentile. Dashed vertical lines indicate timing of debit card receipt.



Figure V: Effect of Debit Cards on Stock of Savings

Sources: Administrative data from Bansefi on account balances and timing of card receipt. Notes: (a) N = 2,023,862 account-period observations from 171,441 accounts over 12 periods. (b) N = 3,086,749 account-period observations from 270,046 accounts over 12 periods. This figure plots ϕ_k from (1). Average balance over each four-month period is the dependent variable, and is winsorized at the 95th percentile. Standard errors are clustered at the bank branch level. Whiskers denote 95% confidence intervals. Black filled in circles indicate results that are significant at the 5% level, gray filled in circles at the 10% level, and hollow circles indicate results that are statistically insignificant from 0. The period prior to receiving the card is the omitted period, which is why its point estimate is 0 with no confidence interval. Dashed vertical lines indicate timing of debit card receipt.



Figure VI: Effect of Debit Cards on Stock of Savings by Distance to Wave 1 Localities

Sources: Administrative data from Bansefi on monthly average account balances and timing of card receipt. Notes: This figure plots ϕ_k from (1), from separate regressions for wave 2 and control beneficiaries closer to a wave 1 locality (orange squares) and farther from a wave 1 locality (black circles). Closer is defined based on a split at the 80th percentile of road distances, based on the test described in the main text. For closer beneficiaries, N =2,294,584 account-period observations from 200,299 accounts over 12 periods. For farther beneficiaries, N =585,085 account-period observations from 50,075 accounts over 12 periods. For farther beneficiaries, N =585,085 account-period observations from 50,075 accounts over 12 periods. Average balance over each four-month period is the dependent variable, and is winsorized at the 95th percentile. Standard errors are clustered at the bank branch level. Whiskers denote 95% confidence intervals. Black filled in circles or orange filled in squares indicate results that are significant at the 5% level, gray filled in circles or light orange filled in squares at the 10% level, and hollow circles or squares indicate results that are statistically insignificant from 0. The period prior to receiving the card is the omitted period, which is why its point estimate is 0 with no confidence interval. Dashed vertical line indicates timing of debit card receipt.



Figure VII: Distribution of Withdrawals per Bimester

Sources: Administrative data from Bansefi on transactions and timing of card receipt. Notes: Based on N = 16,787,160 transactions from 343,204 accounts over 4 years. This figure shows the distribution of withdrawals per bimester from the accounts, before and after the switch to debit cards, for each of wave 1 and wave 2, compared to the control group. The "before wave 1 switches" and "before wave 2 switches" graphs pool all account-bimester pairs before wave 1 or 2 switch to debit cards; similarly, the "after wave 1 switches" and "after wave 2 switches" graphs pool all account-bimester pairs after wave 1 or 2 switch to debit cards.



Figure VIII: Withdrawal to Deposit Ratio Over Time

Sources: Administrative data from Bansefi on transactions and timing of card receipt.

Notes: Based on N = 16,787,160 transactions from 343,204 accounts over 4 years. This figure shows that the withdrawal to deposit ratio is very close to one before receiving the debit card: beneficiaries get a bimonthly deposit from Oportunidades, which they withdraw with one transaction. Immediately after receipt of the card, beneficiaries increase their number of withdrawals, which stays fairly constant thereafter. Dashed vertical lines indicate timing of debit card receipt.



Figure IX: Distribution of Client Deposits per Bimester

Sources: Administrative data from Bansefi on transactions and timing of card receipt. Notes: Based on N = 16,787,160 transactions from 343,204 accounts over 4 years. This figure shows the distribution of client deposits in the accounts, before and after the switch to debit cards, for each of wave 1 and wave 2, compared to the control group. The "before wave 1 switches" and "before wave 2 switches" graphs pool all account-bimester pairs before wave 1 or 2 switch to debit cards; similarly, the "after wave 1 switches" and "after wave 2 switches" graphs pool all account-bimester pairs after wave 1 or 2 switch to debit cards.



Figure X: Effect of Debit Cards on Savings Rate (as Proportion of Income), Wave 1 vs. Control

Sources: Administrative data from Bansefi on account balances by bimester, transactions, and timing of card receipt. Notes: N = 1,852,416 account-period observations from 171,441 accounts over 11 periods. This figure plots $\hat{\Phi}_k$ from (4). Panel (a) is from (3) estimated by Blundell and Bond (1998) two-step system GMM, while panel (b) is from (3) but replacing account fixed effects with a treatment dummy, estimated using OLS. Net balances and transfer amounts are winsorized at the 95th percentile. The variance of $\hat{\Phi}_k$ is estimated using the delta method. Standard errors are clustered at the bank branch level. Whiskers denote 95% confidence intervals. Black filled in circles indicate results that are significant at the 5% level, gray filled in circles at the 10% level, and hollow circles indicate results that are statistically insignificant from 0. The period prior to receiving the card is the omitted period, which is why its point estimate is 0 with no confidence interval. Dashed vertical lines indicate timing of debit card receipt.





Source: Administrative transactions data from Bansefi.

Balance checks

Notes: Based on N = 3,866,525 transactions in 73,070 wave 1 accounts over 4 years. This figure plots the number of balance checks per account tied to a debit card. The black dots shows the total number of balance checks. The gray diamonds show the subset of balance checks which occur after the transfer was received in a bimester, and on a different day than a withdrawal. This subset consists of balance checks that we posit occur to verify if existing balances remain in the account, instead of checking if a new transfer has arrived. Balance checks are zero prior to receiving the card, since it was only possible to check balances at Bansefi branches, which are not charged a fee and hence not recorded in our transactions data. Standard errors are clustered at the bank branch level. Whiskers denote 95% confidence intervals. Dashed vertical line indicates timing of debit card receipt.



Figure XII: Self-Reported Balance Checks and Knowledge

Source: Payment Methods Survey 2012.

Notes: N = 1,617, or less in some regressions if there were respondents who reported "don't know" or refused to respond. Balance checks are measured over the past bimester. Standard errors are clustered at the locality level, using pre-treatment (2004) locality. Whiskers denote 95% confidence intervals. * indicates statistical significance of the difference between those with the card for less vs. more than the median time at p < 0.1, ** p < 0.05, and *** p < 0.01.





Source: Administrative data from Bansefi on transactions and average balances.

Notes: N = 803,770 account-bimester observations from 73,070 wave 1 accounts in 11 bimesters after receiving cards. This figure plots η_c from (6). These coefficients show the within-account net balance difference in pesos, relative to zero balance checks. Standard errors are clustered at the bank branch level. Whiskers denote 95% confidence intervals. Black filled in circles indicate results that are significant at the 5% level, gray filled in circles at the 10% level, and hollow circles indicate results that are statistically insignificant from 0. This figure shows that net balances are significantly lower when beneficiaries check balances more than once per bimester and the difference increases in the number of balance checks, providing support that balance checks are used to monitor the account and build trust.



Figure XIV: Self-Reported Reasons for Not Saving in Bansefi Account

Source: ENCASDU 2010.

Notes: N = 1,694. Standard errors are clustered at the locality level, using pre-treatment (2004) locality. Whiskers denote 95% confidence intervals. * indicates statistical significance of the difference between those with the card for less vs. more than the median time at p < 0.1, ** p < 0.05, and *** p < 0.01.





Source: Administrative transactions data from Bansefi.

Notes: 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 adopt the new technology and use their cards to withdraw their transfers, instead of going to the Bansefi bank branch. Standard errors are clustered at the bank branch level. Whiskers denote 95% confidence intervals. Dashed vertical line indicates timing of debit card receipt.





Sources: ENCELURB panel survey combined with administrative data on timing of card receipt and transfer payment histories for each surveyed beneficiary household.

Notes: N = 9,496 (number of households = 2,942). Dependent variables are measured in pesos per month, with the exception of the asset index. Asset index is the first principal component of assets that are included in both the early (2002, 2003, 2004) and post-treatment (2009–2010) versions of the survey: car, truck, motorcycle, television, video or DVD player, radio or stereo, washer, gas stove, and refrigerator. Standard errors are clustered at the locality level, using pre-treatment (2004) locality. Whiskers denote 95% confidence intervals. Black filled in circles indicate results that are significant at the 5% level, gray filled in circles at the 10% level, and hollow circles indicate results that are statistically insignificant from 0. The * linking consumption and income denotes that a test of equal coefficients from the consumption and income regressions is rejected at the 10 percent level using a stacked regression. Results are from the preferred specification of winsorizing variables at the 95th percentile (and 5th percentile for variables that do not have a lower bound of 0). Raw results, winsorized at 1%, winsorized at 5%, and winsorized at 5% with baseline household characteristics interacted with time fixed effects are available in Appendix Table B.2. All regressions include household and time fixed effects, and standard errors are clustered at the locality level, using pre-treatment (2004) locality.

Study	Intervention	Country	Months		Effect Size		
Drexler, Fischer, and Schoar, 2014	Financial education	Dominican Republic	12	⊢			
Karlan and Zinman, 2016	Interest rate	Philippines	12	\$			
Kast, Meier, and Pomeranz, 2016	Savings group	Chile	12	•			
Ashraf, Karlan, and Yin, 2006	Deposit collection	Philippines	32	- H O H			
Karlan et al., 2016	Reminders	Philippines	9–12	H O H			
Beaman, Karlan, and Thuysbaert, 2014	Savings group	Mali	36				
Dupas et al., 2016	Account	Malawi	24	H O -I			
Callen et al., 2014	Deposit collection	Sri Lanka	25	⊢●1			
Schaner, 2016	Interest rate	Kenya	36				
Kast and Pomeranz, 2014	Account	Chile	12	⊢●-1			
Dupas et al., 2016	Account	Uganda	24				
Dupas and Robinson, 2013	Account or lockbox	Kenya	12	⊢●⊣			
Somville and Vandewalle, 2016	Payment default	India	8	— •			
Prina, 2015	Account	Nepal	13	⊢●-I	1		
Suri and Jack, 2016 (male headed)	Mobile money	Kenya	72	⊢	-		
Brune et al., 2016	Account, commitment	Malawi	7		⊢		
Sayinzoga, Bulte, and Lensink, 2016	Financial education	Rwanda	16	H			
This paper	Debit card	Mexico	20		⊢		
Suri and Jack, 2016 (female headed)	Mobile money	Kenya	72	L	•		
Seshan and Yang, 2014	Financial education	India (migrants to Qatar)	13–17	<			>
			0	025 0.000	0.025	0.050	0.075
			-0.	020 0.000	0.025	0.050	0.075
					Savings Rate		

Figure XVII: Comparison with Other Studies

Notes: For details on how we obtained the effect of savings interventions on the savings rate in each of these studies, as well as additional details about the studies, see Appendix E. Whiskers denote 95% confidence intervals. Black filled in circles indicate results that are significant at the 5% level, gray filled in circles at the 10% level, and hollow circles indicate results that are statistically insignificant from 0. The longer-term (20-month) estimate from this study is represented by an orange square.

Variable	Control	Wave 1	Wave 2	Diff. W1–C	Diff. W2–C	F-test p-value
Panel A: Locality-level data						Г
Log population	10.57	11.18	11.48	0.60***	0.91^{***}	0.000***
	(0.11)	(0.10)	(0.16)	(0.14)	(0.19)	0.000
Bansefi branches per 100,000	1.27	1.23	1.58	-0.03^{-1}	0.32	0.411
1 ,	(0.28)	(0.13)	(0.23)	(0.30)	(0.36)	
% HHs in poverty	$15.93^{'}$	13.20^{-1}	$12.23^{'}$	-2.73°	-3.71^{*}	0.177
	(1.67)	(0.75)	(1.09)	(1.82)	(1.99)	
Occupants per room	1.18	1.11	1.12	-0.07	-0.06	0.260
	(0.04)	(0.01)	(0.02)	(0.04)	(0.04)	
Number of localities	44	143	88			
Panel B: Administrative bank acco	unt data					
Average balance	581.25	670.32	614.29	89.07	33.05	0.112
5	(12.46)	(56.24)	(21.26)	(55.33)	(23.95)	
Number of deposits	1.00	0.99	1.00	-0.01^{**}	0.00	0.040^{**}
	(0.00)	(0.00)	(0.01)	(0.00)	(0.01)	
Number of withdrawals per deposi	t 0.96	0.96	0.96	0.00	-0.01	0.811
	(0.00)	(0.02)	(0.01)	(0.02)	(0.01)	
Percent withdrawn	99.22	99.31	99.37	0.08	0.15	0.703
	(0.10)	(0.28)	(0.16)	(0.30)	(0.18)	
Size of transfer	1292.39	1507.09	1523.55	214.70^{***}	231.16^{***}	0.000^{***}
	(10.86)	(18.76)	(13.85)	(21.27)	(16.98)	
Years with account	5.51	5.58	5.73	0.06	0.22	0.554
	(0.08)	(0.14)	(0.20)	(0.16)	(0.21)	
Number of accounts	$97,\!918$	$73,\!070$	171,717			

Table I: Comparison of Baseline Means

Sources: Census (2005), Bansefi branch locations (2008), poverty estimates from Oportunidades (based on 2005 Census), and administrative data from Bansefi.

Notes: HHs = households, W1 = wave 1, W2 = wave 2, C = control, Diff. = difference. For the administrative data from Bansefi, baseline is defined as November 2007 to October 2009 (prior to any accounts receiving cards in the data from Bansefi).

	(1)	(2)	(3)
	OLS	2SLS	2SLS
Coefficient	0.001	0.028**	0.029^{*}
	(0.002)	(0.013)	(0.014)
First stage F-test for $Trust_{it}$		40.0	18.1
First stage F-test for $Trust_{it} \times Net \ Balance_{i,t-1}$			147.3
First stage F-test for $Trust_{it} \times Transfers_{it}$			38.3
Number of observations	1330	1330	1330
Lagged balance and transfers	No	No	Yes

Table II: Relationship between Trust and Savings Rates

Sources: ENCASDU survey data merged with administrative bank account balance and transactions data from Bansefi.

Notes: N = 1,330 beneficiary households merged with accounts. The specification for column 1 is $\hat{\zeta}/\overline{Y}$ from (8) with OLS; column 2 is $\hat{\zeta}/\overline{Y}$ from (8) with 2SLS, instrumenting trust with a set of dummies for timing of card receipt; column 3 is $\hat{\Phi} = (\hat{\zeta} + \hat{\xi}\omega_{-1} + \hat{\psi}\mu)/\overline{Y}$ from

$$\begin{split} Net \ Balance_{it} - Net \ Balance_{i,t-1} &= \zeta Trust_{it} + \theta Net \ Balance_{i,t-1} + \xi Trust_{it} \times Net \ Balance_{i,t-1} \\ &+ \gamma Transfers_{it} + \psi Trust_{it} \times Transfers_{it} + \varepsilon_{it} \end{split}$$

with 2SLS, instrumenting trust and its interactions with lagged net balance and transfers with a set of dummies for timing of card receipt and their interactions with lagged net balance and transfers. Coefficients are expressed as a proportion of average income.

Supplementary Material (FOR ONLINE PUBLICATION ONLY)

Appendix A Sample of Materials Received by Beneficiaries



Figure A.1: Flyer Provided with the Debit Card (Front)

Notes: This flyer is provided by Oportunidades together with the debit card. The front of the flyer provides activation instructions and security tips regarding the PIN number and debit card.



Figure A.2: Flyer Provided with the Debit Card (Back)

Notes: The back of the flyer provides instructions on using the card to withdraw money at ATMs and to make purchases. It clarifies that the card can be used to withdraw money at any ATM within the networks RED and PLUS (which cover almost all ATMs in Mexico) and at major grocery stores. Figure A.3: Sample Calendar of Transfer Dates Given to Beneficiaries

Oport <u>unidades</u>	Calendario F de Apoyos	ijo de Retiro SEDESOL		
Entidad: 15 MEXICO		Folio Titular:		
Zona de Atención: 150	303	Nombre Titular:		
Municipio: 33 ECATER	PEC DE MORELOS	Identificador de Familia:		
Localidad: 1 ECATEP	EC DE MORELOS	Fase de Incorporación: 35		
AGEB: Código Postal: 55450		Esquema de Apoyos: Urbano 1		
Domicilio:		Colonia:		
Domicilio:	os del bimestre de corresponsabili	Colonia:		
Domicilio: Estimada Titular: Los apoy	os del bimestre de corresponsabili viembre - Diciembre del 2008	Colonia: dad los puede retirar a partir del Lunes 20 de Abril del 2009		
Domicilio: Estimada Titular: Los apoye	os del bimestre de corresponsabili viembre - Diciembre del 2008 Enero - Febrero del 2009	Colonia: dad los puede retirar a partir del Lunes 20 de Abril del 2009 Lunes 1 de Junio del 2009		
Domicilio:	os del bimestre de corresponsabili riembre - Diciembre del 2008 Enero - Febrero del 2009 Marzo - Abril del 2009	Colonia: dad tos puede retirar a partir del Lunes 20 de Abril del 2009 Lunes 1 de Junio del 2009 Lunes 13 de Julio del 2009		
Domicilio:	os del bimestre de corresponsabili viembre - Diciembre del 2008 Enero - Febrero del 2009 Marzo - Abril del 2009 Mayo - Junío del 2009	Colonia: dad tos puede retirar a partir del Lunes 20 de Abril del 2009 Lunes 1 de Junio del 2009 Lunes 13 de Julio del 2009 Lunes 14 de Septiembre del 2009		
Domicilio:	os del bimestre de corresponsabili viembre - Diciembre del 2008 Enero - Febrero del 2009 Marzo - Abril del 2009 Mayo - Junio del 2009 Julio - Agosto del 2009	Colonia: dad Ios puede retirar a partir del Lunes 20 de Abril del 2009 Lunes 1 de Junio del 2009 Lunes 13 de Julio del 2009 Lunes 14 de Septiembre del 2009 Lunes 16 de Noviembre del 2009		

Titular beneficiaria: Usted podrá retirar sus apoyos con su Tarjeta de Débito a partir de la fecha indicada en cajeros automáticos ó establecimientos autorizados (que aceptan tarjetas VISA).

Recuerde que en cajeros automáticos podrá realizar dos operaciones (retiros ó consultas) gratuitas al bimestre, también puede utilizar su Tarjeta para comprar en establecimientos que aceptan Tarjetas de Débito VISA.

Para mayor Información, consultas, dudas ó quejas, comunicarse al 01800 500 50 50 de lunes a viernes de 9 de la mañana a 6 de la tarde.

"Este programa es público, ajeno a cualquier partido político. Queda prohibido el uso para fines distintos al desarrollo social."

"Este programa es de carácter público, no es patrocinado ni promovido por partido político alguno y sus recursos provienen de los impuestos que pagan todos los contribuyentes. Está prohibido el uso de este programa con fines políticos, electorales, de lucro y de otros distintos a los establecidos. Quien haga uso indebido de los recursos de este programa deberá ser denunciado y sancionado de acuerdo a la ley aplicable y ante la autoridad competente "

Consecutivo: 1005

Notes: This is a sample of the calendars that provide the transfer dates to recipients. For each bimester in the following year, it states the corresponding payment date. It reminds recipients that they should use their debit cards after the indicated date at ATMs or establishments accepting VISA. It also reminds them that they are allowed two free transactions per bimester at ATMs.

Appendix B Additional Figures and Tables



Figure B.1: Geographic Coverage and Expansion of Debit Cards

Sources: Administrative data from Oportunidades on timing of debit card receipt by locality and shape files from INEGI.

Notes: N = 275 localities (44 in control, 143 in wave 1, 88 in wave 2). The area of each urban locality included in the study is shaded according to its wave of treatment. Urban localities that were not included in the Oportunidades program at baseline or were included in the program but did not pay beneficiaries through Bansefi savings accounts are not included in the figure or in our study.



Figure B.2: Effect of Debit Cards on Stock of Savings (Net Balance)

Source: Administrative data from Bansefi on transactions and average balances. Notes: (a) N = 2,023,862 from 171,441 accounts over 12 periods. (b) N = 3,086,749 from 270,046 accounts over 12 periods. The figure plots ϕ_k from (1) with Net Balance_{it} as the dependent variable. Net balances refer to average balances minus the mechanical effect on average balance of leaving a portion of the deposit in the account for a certain number of days before withdrawing it during the same bimester. Average balance over each four-month period is the dependent variable, and is winsorized at the 95th percentile. Whiskers denote 95% confidence intervals. Black filled in circles indicate results that are significant at the 5% level, gray filled in circles at the 10% level, and hollow circles indicate results that are statistically insignificant from 0. The period prior to receiving the card is the omitted period, which is why its point estimate is 0 with no confidence interval. Dashed vertical lines indicate timing of debit card receipt.



Figure B.3: Effect of Debit Card on Savings Rate without Transfer Interactions

Source: Administrative data from Bansefi on transactions and average balances. Notes: N = 1,852,416 account-period observations from 171,441 accounts over 11 periods. This figure plots $\hat{\Phi}_k \equiv (\hat{\alpha}_k + \hat{\xi}_k \omega_{k-1})/\overline{Y}$ from

$$\begin{split} &Net \; Balance_{it} - Net \; Balance_{i,t-1} = \\ &\lambda_i + \sum_k \delta_k \mathbb{I}(k=t) + \sum_k \alpha_k T_{j(i)} \times \mathbb{I}(k=t) \\ &+ \sum_k \theta_k Net \; Balance_{i,t-1} \times \mathbb{I}(k=t) + \sum_k \xi_k Net \; Balance_{i,t-1} \times T_{j(i)} \times \mathbb{I}(k=t) + \varepsilon_{it}. \end{split}$$

Panel (a) is estimated by Blundell and Bond (1998) two-step system GMM, while panel (b) is from the above specification but replacing account fixed effects with a treatment dummy, estimated using OLS. Net balances are winsorized at the 95th percentile. The variance of $\hat{\Phi}_k$ is estimated using the delta method. Whiskers denote 95% confidence intervals. Black filled in circles indicate results that are significant at the 5% level, gray filled in circles at the 10% level, and hollow circles indicate results that are statistically insignificant from 0. The period prior to receiving the card is the omitted period, which is why its point estimate is 0 with no confidence interval. Dashed vertical lines indicate timing of debit card receipt.



Figure B.4: Separated by Time with Account: Effect of Debit Cards on Savings Rate (as Proportion of Income), Wave 1 vs. Control

Source: Administrative data from Bansefi on transactions and average balances.

Notes: (a) and (b) N = 792,912 from 75,000 accounts; (c) and (d) N = 1,055,890 from 95,990 accounts. This figure plots $\hat{\Phi}_k$ from (4), separately for older and younger accounts. Accounts are split into older accounts and younger accounts based on the median account opening date, which is October 19, 2004. Panels (a) and (c) are from (3) estimated by Blundell and Bond (1998) two-step system GMM, while (b) and (d) are from (3) but replacing account fixed effects with a treatment dummy, estimated using OLS. Net balances and transfer amounts are winsorized at the 95th percentile. The variance of $\hat{\Phi}_k$ is estimated using the delta method. Standard errors are clustered at the bank branch level. Whiskers denote 95% confidence intervals. Black filled in circles indicate results that are significant at the 5% level, gray filled in circles at the 10% level, and hollow circles indicate results that are statistically insignificant from 0. The period prior to receiving the card is the omitted period, which is why its point estimate is 0 with no confidence interval. Dashed vertical lines indicate timing of debit card receipt.

	(1)	(2)	(3)
	Mean for Card	Difference for Card	P-value of
	> Median Time	< Median Time	Difference
Number of household members	5.18	0.26	0.114
	(0.08)	(0.15)	
Number of children	2.19	0.03	0.743
	(0.08)	(0.10)	
Age of household head	44.73	0.96	0.246
	(0.08)	(0.80)	
Household head is male	0.67	0.02	0.603
	(0.03)	(0.03)	
Household head is married	0.70	0.02	0.459
	(0.04)	(0.03)	
Education level of head	9.30	-0.33	0.092^{*}
	(0.16)	(0.18)	
Occupants per room	3.50	-0.03	0.801
	(0.07)	(0.11)	
Access to health insurance	0.59	0.05	0.165
	(0.02)	(0.03)	
Asset index	0.04	-0.04	0.605
	(0.04)	(0.08)	
Income	3190.32	222.69	0.150
	(47.40)	(146.67)	

Table B.1: Balance test in ENCASDU

Source: ENCASDU 2010.

Notes: N = 1,694, or less for variables that were missing for some observations. Standard errors clustered at the locality level.

	(1)	(2)	(3)	(4)	Mean
Consumption	-178.11^{**}	-153.96^{**}	-138.09^{**}	-143.63^{**}	2731.20
	(80.15)	(69.49)	(60.86)	(62.11)	(82.81)
Income	78.98	85.09	49.44	46.28	3148.28
	(168.11)	(149.46)	(128.00)	(130.40)	(89.02)
P-value Consumption vs. Income	$[0.058]^*$	$[0.055]^*$	$[0.092]^*$	[0.103]	
Savings = Income - Consumption	257.09^{*}	243.20**	236.16^{**}	243.75**	412.17
	(132.50)	(118.50)	(102.04)	(108.26)	(103.32)
Purchase of durables	9.77	8.64	8.20	7.54	32.98
	(12.41)	(8.61)	(4.99)	(4.98)	(3.32)
Asset index	0.06	0.06	0.08	0.07	0.48
	(0.08)	(0.08)	(0.07)	(0.07)	(0.10)
Number of households	$2,\!942$	$2,\!942$	$2,\!942$	$2,\!929$	
Number of observations	$9,\!496$	$9,\!496$	$9,\!496$	$9,\!469$	
Time fixed effects	Yes	Yes	Yes	Yes	
Household fixed effects	Yes	Yes	Yes	Yes	
Household characteristics \times time	No	No	No	Yes	
Winsorized	No	1%	5%	5%	5%

Table B.2: Change in Savings and Assets After Receiving Card

Sources: ENCELURB panel data merged with administrative data on beneficiary status and timing of debit card receipt.

Notes: Each row label is the dependent variable from a separate regression; each column is a different specification. The "Mean" column shows the mean of the dependent variable for the control group, winsorized at 5%. * indicates statistical significance at p < 0.1, ** p < 0.05, and *** p < 0.01. Standard errors are clustered at the locality level, using pre-treatment (2004) locality. Dependent variables are measured in pesos per month, with the exception of the asset index. Asset index is the first principal component of assets that are included in both the early (2002, 2003, 2004) and post-treatment (2009-2010) versions of the survey: car, truck, motorcycle, television, video or DVD player, radio or stereo, washer, gas stove, and refrigerator. Household characteristics are measured at baseline (2004, or for households that were not included in the 2004 wave, 2003). They include characteristics of the household head (working status, a quadratic polynomial in years of schooling, and a quadratic polynomial in age), whether anyone in the household has a bank account, a number of characteristics used by the Mexican government to target social programs (the proportion of household members with access to health insurance, the proportion age 15 and older that are illiterate, the proportion ages 6-14 that do not attend school, the proportion 15 and older with incomplete primary education, the proportion ages 15-29 with less than 9 years of schooling), and dwelling characteristics (dirt floors, no bathroom, no piped water, no sewage, and number of occupants per room). The number of households in column (4) is slightly lower because 13 households have missing values for one of the household characteristics included (interacted with time fixed effects) in that specification.

	Т	Total		nsefi
	ATMs	Branches	ATMs	Branches
Current quarter	-0.37	-0.01	0.00	-0.01
	(1.51)	(0.34)	(0.00)	(0.02)
1 quarter lag	-1.79	0.10	-0.01	0.02
	(2.49)	(0.37)	(0.01)	(0.02)
2 quarter lag	2.04	0.12	0.01	0.01
	(3.72)	(0.39)	(0.01)	(0.02)
3 quarter lag	-0.57	-0.01	-0.01	0.02
	(1.11)	(0.29)	(0.01)	(0.02)
4 quarter lag	2.29	-0.28	0.00	-0.04
	(2.54)	(0.64)	(0.00)	(0.03)
5 quarter lag	-1.13	0.08	0.00	0.00
	(2.56)	(0.81)	(0.00)	(0.02)
6 quarter lag	-0.31	0.94	0.00	0.02
	(3.60)	(0.67)	(0.00)	(0.02)
1 quarter lead	0.66	-0.25	0.00	-0.01
	(1.74)	(0.40)	(0.00)	(0.02)
2 quarter lead	3.96	0.11	0.01	0.00
	(3.65)	(0.40)	(0.01)	(0.02)
3 quarter lead	-0.06	0.26	-0.01	-0.01
	(4.18)	(0.65)	(0.02)	(0.03)
4 quarter lead	-2.50	0.83	0.00	-0.04
	(4.04)	(0.78)	(0.01)	(0.05)
5 quarter lead	3.97	0.27	0.00	0.01
	(3.19)	(0.40)	(0.00)	(0.02)
6 quarter lead	5.18^{*}	-0.98	0.01	-0.04
	(3.03)	(0.97)	(0.01)	(0.03)
Mean control group	46.08	37.13	0.09	1.42
F-test of lags	0.59	0.60	0.73	1.15
[p-value]	[0.74]	[0.73]	[0.63]	[0.33]
F-test of leads	0.87	1.00	1.24	0.79
[p-value]	[0.52]	[0.42]	[0.29]	[0.58]
Municipality fixed effects	Yes	Yes	Yes	Yes
Quarter fixed effects	Yes	Yes	Yes	Yes

Table B.3: Supply-Side Response

Source: Data obtained from CNBV.

Notes: N = 2,491 municipality-quarter observations from 199 municipalities. This table shows β_k from (10). The F-test of lags tests $\beta_{-6} = \cdots = \beta_{-1} = 0$; the F-test of leads tests $\beta_1 = \cdots = \beta_6 = 0$. * indicates statistical significance at p < 0.1, ** p < 0.05, and *** p < 0.01.

Dependent variable	p-value
Consumption	0.322
Income	0.159
Savings = Income - Consumption	0.176
Purchase of Durables	0.269
Asset Index	0.398
Number of households	$2,\!942$
Number of observations	$9,\!496$
Household fixed effects	Yes
Time fixed effects	Yes
Winsorized	5%

Table B.4: Parallel Trends in Consumption, Income, Savings, Assets

Sources: ENCELURB panel data merged with administrative data on beneficiary status and timing of debit card receipt.

Notes: This table shows p-values from an F-test of $\omega_k = 0 \forall k < 2009$ (where k = 2002 is the reference period and is thus omitted) from

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

Dependent variables are measured in pesos per month, with the exception of the asset index. Asset index is the first principal component of assets that are included in both the early (2002, 2003, 2004) and post-treatment (2009–2010) versions of the survey: car, truck, motorcycle, television, video or DVD player, radio or stereo, washer, gas stove, and refrigerator. Dependent variables are winsorized at the 5% level.

Appendix C Mechanical Effect

This appendix defines the "mechanical effect," which we use to compute net balances. We explain the logic behind the mechanical effect, present an example, and provide a step by step guide for its computation, summarized in Table C.1.

C.1 Logic of the Mechanical Effect

The mechanical effect is the contribution to average balances from the transit of transfers in recipients' accounts. Since the mechanical effect does not represent net (long-term) savings, or even saving from one period to the next, our goal is to net it out from average balances and construct a measure of net balances, *Net Balance_{it}*. Changes in the mechanical effect can arise due to changes in the *frequency* of withdrawals. For example, if client A begins the period with 0 balance, receives 2,000 pesos in her account, and withdraws 1,000 pesos on the first day of the period, and the other 1,000 pesos midway through the period, her average balance will equal 1,000 * 0 + 1,000 * $\frac{1}{2} = 500$ pesos. Compared to client B who withdrew the entire 2,000 pesos on the first day of the period, client A's average balance is 500 pesos higher, but both end the period with a balance of zero. Their net balances, constructed as average balance minus mechanical effect, are both equal to zero.

Changes in the mechanical effect can also arise from changes in the *timing* of withdrawals, compared to the deposit dates. The deposit date is usually known by the recipients: Oportunidades generally disburses transfers within the first week of the bimester, and the program distributes calendars stating the dates when accounts will be credited. Nevertheless, beneficiaries may not withdraw their benefits on the day they are deposited, which also leads to a mechanical effect that contributes to the average balance. In our data, the mechanical effect can thus change for debit card recipients relative to the control group as a result of increased withdrawal frequency of smaller amounts and changes in time between the deposit and first withdrawal.

Finally, we need to compare not only the timing of deposits and withdrawals, but also their relative sizes. Although the calculation is simple, there are several cases to consider depending on the number of withdrawals, when they occur, and whether they exceed the amount deposited that period. We use an example to exemplify the steps involved.

C.2 Example:

- Select a pattern where clients received a single deposit (the most common, although as explained previously, beneficiaries receive more than one Oportunidades deposit in some bimesters)
- 2. Select a pattern with one deposit followed by two withdrawals (DWW)
- 3. The pattern with one deposit and two withdrawals (DWW), must fit in one of the following three scenarios: (a) the deposit is less than the first withdrawal $(W_1 \ge D)$, (b) the deposit is larger than the first withdrawal but smaller than the sum of the two withdrawals $(W_1 < D \& W_1 + W_2 \ge D)$, (c) the deposit is larger than the sum of withdrawals $(W_1 + W_2 < D)$.

- 4. Compute the mechanical effect, at the individual level, for each of the three scenarios discussed above:
 - (a) The deposit is less than the first withdrawal \Rightarrow the mechanical effect is just the time lapse between the deposit and the first withdrawal times the deposit amount $(lapse_{DW_1} * D)$.
 - (b) The deposit is larger than the first withdrawal but smaller than the sum of the two withdrawals \Rightarrow the mechanical effect is the time lapse between the deposit and the first withdrawal times the amount of the first withdrawal, plus the time lapse between the deposit and the second withdrawal times the remaining deposit amount after subtracting the first withdrawal ($lapse_{DW_1} * W_1 + lapse_{DW_2} * (D W_1)$).
 - (c) The deposit is larger than the sum of the withdrawals \Rightarrow the mechanical effect is the time lapse between the deposit and the first withdrawal times the amount of the first withdrawal, plus the time lapse between the deposit and the second withdrawal times the amount of the second withdrawal ($lapse_{DW_1} * W_1 + lapse_{DW_2} * (W_2)$).

Table C.1 shows the cases we considered as well as their prevalence in the data.

C.3 Steps

More generally we follow the steps below:

- 1. We separate the sample based on the number of transfers received by Opportunidades' beneficiaries: 85% of beneficiary-bimester pairs receive a single transfer in the bimester and 15% received two transfers in the same bimester. See footnote 18 for a description of the reasons some beneficiary-bimester pairs include more than one transfer.
- 2. We determine the pattern of transactions: for example, a beneficiary who first received a deposit and then performed two withdrawals has a sequence (D, W_1, W_2) , or DWW for short.
- 3. We compare the size of the deposit to the withdrawals, and generate different scenarios. These scenarios depend on the relative size of the deposit and withdrawals: each withdrawal could be larger than the deposit, their sum might be larger, or the deposit is larger than the sum of withdrawals.
- 4. We compute the mechanical effect. To do this, we measure the lapse of time, in days, which passes between the deposit and each withdrawal, and multiply the time lapses by the amount of the transfer which only transited through the account, and was not kept in the account through the end of and into the next bimester.

	Pattern	% Total	Conditions	Mechanical Effect
Pan (1)	el A. Regul DW	ar patterns 73.4	: single deposit into account in the bime $W \le D$ W > D	ester $lapse_{DW} * W$ $lapse_{DW} * D$
(2)	DWW	9.1	$W_1 \ge D$ $W_1 < D \& W_1 + W_2 \ge D$ $W_1 + W_2 < D$	$lapse_{DW_1} * D$ $lapse_{DW_1} * W_1 + lapse_{DW_2} * (D - W_1)$ $lapse_{DW_1} * W_1 + lapse_{DW_2} * (W_2)$
(3)	DWWW	1.7	$W_1 \ge D W_1 < D \& W_1 + W_2 \ge D W_1 + W_2 < D \& W_1 + W_2 + W_3 \ge D$	$\begin{split} &lapse_{DW_{1}}*D\\ &lapse_{DW_{1}}*W_{1}+lapse_{DW_{2}}*(D-W_{1})\\ &lapse_{DW_{1}}*W_{1}+lapse_{DW_{2}}*W_{2}\\ &+lapse_{DW_{3}}*(D-W_{1}-W_{2}) \end{split}$
Pan	el B. Irregu	lar pattern	es: multiple deposits into account in the	bimester
(4)	DDWW	3.1	$ \begin{split} W_1 &\leq D_1 \& W_2 \leq D_2 \\ W_1 &> D_1 \& W_2 \leq D_2 \\ W_1 &\leq D_1 \& W_2 < D_2 \\ W_1 &> D_1 \& W_2 > D_2 \end{split} $	$\begin{split} lapse_{D_1W_1} * W_1 + lapse_{D_2W_2} * W_2 \\ lapse_{D_1W_1} * D_1 + lapse_{D_2W_2} * W_2 \\ lapse_{D_1W_1} * W_1 + lapse_{D_2W_2} * D_2 \\ lapse_{D_1W_1} * D_1 + lapse_{D_2W_2} * D_2 \end{split}$
(5)	DWD	3.0	$W \le D_1$ $W > D_1$	$lapse_{D_1W} * W$ $lapse_{D_1W} * D_1$
(6)	DDW	2.7	$\begin{split} W &\geq D_1 + D_2 \\ W &< D_1 + D_2 \& W \leq D_2 \\ W &< D_2 \end{split}$	$lapse_{D_1W} * D_1 + lapse_{D_2W} * D_2$ $lapse_{D_1W} * (W - D_2) + lapse_{D_2W} * D_2$ $lapse_{D_2W} * W$
(7)	DWDW	1.6	$ \begin{split} W_1 &\leq D_1 \& W_2 \leq D_2 \\ W_1 &> D_1 \& W_2 \leq D_2 \\ W_1 &\leq D_1 \& W_2 < D_2 \\ W_1 &> D_1 \& W_2 > D_2 \end{split} $	$\begin{split} lapse_{D_1W_1} * W_1 + lapse_{D_2W_2} * W_2 \\ lapse_{D_1W_1} * D_1 + lapse_{D_2W_2} * W_2 \\ lapse_{D_1W_1} * W_1 + lapse_{D_2W_2} * D_2 \\ lapse_{D_1W_1} * D_1 + lapse_{D_2W_2} * D_2 \end{split}$

Table C.1: Computation of Mechanical Effect

Notes: D_i indicates the *i*th deposit and W_i indicates the *i*th withdrawal within a bimester. $lapse_{D_iW_j}$ measures the number of days between the *i*th deposit and the *j*th withdrawal, divided by the number of days in the bimester. The patterns listed here represent 95% of all bimonthly patterns, but all patterns representing at least 0.01% of all account-bimester pair patterns have been coded to obtain an estimate of the mechanical effect.
Appendix D Details on the GMM estimation

In this section we clarify some details of the GMM estimation. First, standard errors of the parameters in (3) are clustered at the bank branch level and corrected for finite sample bias following Windmeijer (2005); the formula for the variance of Φ_k is then approximated using the delta method.

Because lags of lagged net balance and its interactions are used as GMM-style instruments in the system GMM estimator, including Net Balance_{i,t-1}× $\mathbb{I}(k = t)$ and Net Balance_{i,t-1}× $T_{j(i)}$ × $\mathbb{I}(k = t)$ terms for all k leads to a proliferation of instruments and an over-parameterized model. As a result, for the lagged balance terms we create three groups of time dummies (focusing on wave 1, since we do not have enough post-treatment periods in wave 2 to explore dynamic effects): pre-treatment, post-treatment learning period, and post-learning period, and interact these (rather than $\mathbb{I}(k = t)$ time dummies for each k) with lagged net balance. We thus estimate the dampening effect of lagged net balance between the treatment and control group during the three broad periods (θ_1 , θ_2 , and θ_3), and the difference in the dampening effect of lagged net balance between the treatment and control group during the

The details of our two-step system GMM estimation are as follows. Lagged balance is used as an endogenous GMM-style instrument; because bias can increase in finite samples as T increases (since this leads to more lags and, hence, more instruments: see Ziliak, 1997), to reduce the number of instruments we collapse instruments (as in Beck and Levine, 2004), following the procedure described in Roodman (2009). Because $Transfers_{it}$ is predetermined but not strictly exogenous, variables on the right-hand side of (3) interacted with $Transfers_{it}$ are valid instruments in the system's equation in levels, but not the equation in differences; as a result, we include time dummies and all interaction terms on the right-hand side of (3) as IV-style instruments in the system's equation in levels, and time dummies and interaction terms excluding those interacted with $Transfers_{it}$ in the equation in differences. These specification choices result in a total instrument count of 88. Because our panel does not include gaps, we use first differencing—as in Blundell and Bond (1998)—rather than the sample-maximizing forward orthogonal deviations—as in Arellano and Bover (1995)—to eliminate fixed effects in the transformed equation to be estimated.

Our motivating precautionary saving model makes a number of testable predictions about the dynamic effect of lagged balance on the flow of savings, i.e., the θ_k and ξ_k terms from (3). As described above, including Net Balance_{i,t-1} \times \mathbb{I}(k = t) and Net Balance_{i,t-1} $\times T_{j(i)} \times \mathbb{I}(k = t)$ terms for all k leads to a proliferation of instruments in the GMM estimation, s o we estimate the dampening effect of lagged net balance during three broad periods: pre-treatment, post-treatment pre-learning, and post-learning. For these three broad periods, the coefficients θ_p for $p \in \{1, 2, 3\}$ give the dampening effect for the control group, and ξ_p give the differential dampening effect for the control.

In the control group, the households have not been given a technology to monitor and build trust in the bank, so we expect their precautionary savings target *in the account* to be 0. Thus, if the household's balance increases by a bit in the previous period, the balance would exceed the savings target, and the household would fully offset this increase in Net $Balance_{i,t-1}$ by dissaving in period t, resulting in a strong negative correlation between lagged balance and the flow of savings, and thus $\theta_p \approx -1$ for all p. This prediction is fairly accurate: we find $\hat{\theta}_1$, $\hat{\theta}_2$, and $\hat{\theta}_3$ ranging from -0.76to -0.59. In the treatment group, we expect a similar effect in the pre-treatment period and the learning period since treatment beneficiaries do not yet trust the bank in these periods, i.e. $\xi_1 = 0$, $\xi_2 = 0$; the prediction is again fairly accurate, with $\hat{\xi}_1 = 0.10$, $\hat{\xi}_2 = 0.12$.

Once households have built trust in the bank through balance checks, the savings target in the account would no longer be 0, but some positive number less than or equal to the household's overall precautionary savings target. Once the precautionary savings target in the account has increased, the household would no longer fully offset increases in Net Balance_{i,t-1}, so the dampening effect of net balance would shift from close to -1 toward 0, so (with $\theta_3 \approx -1$) we expect $0 < \xi_3 < 1$; because precautionary savings models predict that the flow of saving is decreasing in lagged assets as the level of precautionary savings approaches its target, we nevertheless expect $\theta_3 + \xi_3 < 0$. This is indeed what we find: $\hat{\xi}_3 = 0.58$, and $\hat{\theta}_3 + \hat{\xi}_3 = -0.18$.

Appendix E Comparison with Other Studies

The savings rates in Figure XVII are drawn form papers which meet the following five criteria.

- We try to include all studies measuring the impact of savings interventions. This includes
 offering accounts or other savings devices, deposit collection, financial education, and savings
 group interventions, as well as sending reminders, changing the interest rate, and defaulting
 payments. We exclude studies which measure the impact of income shocks and cash transfers
 on savings, since these are not savings interventions.
- 2. We only include studies with a duration of at least 6 months.
- 3. We focus on interventions in developing countries aimed at adults.
- 4. We include papers published in peer-reviewed journals, NBER working papers, and other working papers listed as "revise and resubmit" on authors' websites as of January 2017. This filter intends to avoid using preliminary results.
- 5. Finally, to estimate the savings rate we need to divide the change in savings by income. We therefore only include studies that include average income (or, at a minimum, consumption or profits) in their tables, or an income (or consumption or profits) variable in the replication data.

Most papers report the impact of savings interventions on savings balances, which we divide by total income over the relevant period to obtain a savings rate. We use intent-to-treat estimates. In the cases that replication data are available, we use the replication data to replicate the studies' findings and compute the intent-to-treat impact of the intervention on the savings rate. When possible, we use total savings; when this is not available, we use savings in the savings intervention being studies (e.g., in the bank). For studies with results for multiple periods in time, we select results for the longer time period. This appendix provides more detail on how the savings rates in Figure XVII were computed for each study.

Ashraf, Karlan, and Yin (2006a). This study looks at the effect of a deposit collection service in the Philippines. The authors find an effect of the deposit collection service on bank savings after 12 months that is statistically significant at the 10% level, but that dissipates and is no longer significant after 32 months; the effect on total savings after 12 months is of similar magnitude to that of bank savings, but is noisier and not statistically significant. We use the effect on bank savings after 32 months (since the effect on total savings after 32 months is not available). The effect on bank savings after 32 months is 163.52 pesos (Table 6), which we divide by total income over 32 months, which was obtained by dividing annual household income (129,800 pesos; Table 1, column 2 of the December 2005 version but not included in the final version) by 12 months to get monthly income, then multiplying by the 32-month duration of the study.

Beaman, Karlan, and Thuysbaert (2014). This study looks at the effect of introducing rotating savings and credit association (ROSCA) groups in Mali to new techniques in order to improve their flexibility, namely allowing members to take out loans from the group savings rather than waiting for their turn to take home the whole pot. We use the impact of treatment on total savings (Table 4A, column 2) of \$3.654. Because income was not included in their survey, we calculate expenditure over the three year period to use as the denominator by adding monthly food consumption (weekly food consumption from Table 5, column 6, times the average number of weeks per month) and monthly non-food consumption (Table 5, column 5), and multiplying this by the 36 months of the study. Since these consumption figures are expressed in per adult equivalent rather than per household amounts, we then multiply by the average number of adult equivalents—we conservatively estimate this as 1+0.7+ (average number of household members -2), where average number of household members, 7.55, is from Table 3. This estimate uses the OECD adult equivalence scale and assumes that only two household members are adults, thus leading us to likely underestimate its true value and thereby overestimate the savings rate.

Brune et al. (2016). This study looks at the effect of allowing farmers in Malawi to channel profits from their harvests into formal bank accounts; some farmers are also offered a commitment account. We use the intent-to-treat impact of any account on bank savings after 7 months, 1863 kwacha (Table 5, column 1). We divide by average monthly expenditure in the treatment group from Table 6, multiplied by the 7 months in the study.

Callen et al. (2014). This study looks at the effect of offering deposit collection to rural households in Sri Lanka. We use the estimate of the impact of treatment on savings from Table 1 (of the February 2016 version of the paper, which is more recent than the NBER version), or 883 Rupees. Since this effect pools results from surveys conducted 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 19, and 25 months post-treatment, we divide the savings effect by total household income over the previous month from Table A3 times the average number of post-treatment months (9 months).

Drexler, Fischer, and Schoar (2014). This study looks at the effect of financial literacy training in the Dominican Republic. In the study, neither the standard accounting nor rules of thumb treatment arms have a statistically significant impact on savings. We use the replication microdata to replicate their results from Table 2 of the impact of training on savings; we then estimate the pooled treatment effect. Because the paper and data set do not include income or expenditures, we use microenterprise sales in the denominator (the sample consisted entirely of microentrepreneurs). We calculate average sales among the treatment group at endline in the microdata, and multiply this by the 12-month duration of the study.

Dupas and Robinson (2013b). This study looks at the effect of providing different savings tools to ROSCA members in Kenya: a savings box, locked savings box, health savings pot, and health savings account. We used replication data to replicate the results in the paper and estimate a pooled treatment effect for the three interventions in which savings could be directly measured: the savings box, lockbox, and health savings account. We divide the savings effect by average weekly income among the treatment group (which we calculate using the replication data) multiplied by the 52-week duration of the study.

Dupas et al. (2016b). This study looks at the impact of providing access to formal savings accounts to households in three countries: Chile, Malawi, and Uganda. In Chile, an endline survey was not conducted due to low take-up, so we cannot include results for this country. For Malawi and Uganda, we use the intent-to-treat impact of treatment on total monetary savings of \$1.39 in Uganda and \$4.98 in Malawi (Table 5, column 7). Because these effects pool data from three surveys conducted 12, 18, and 24 months post-treatment, we divide the savings effect by monthly household income (Table 10) multiplied by the average post-treatment time (18 months).

Karlan et al. (2016). This study looks at the effect of text message reminders to save in Bolivia, Peru, and the Philippines. Because the Philippines is the only country for which income data was collected, it is the only country from the study for which we estimate the effect of treatment on the savings rate. We use replication data to estimate the effect of treatment on the level of savings. (The paper uses a log specification, but for consistency with the other studies we use levels; in both cases, the effect is statistically insignificant for the Philippines.) Because savings was measured between 9 and 11 months after treatment, we divide by average weekly income of the treatment group (estimated using the replication data) times the average number of weeks per month times the midpoint number of months (10 months).

Karlan and Zinman (2016). This study looks at the effect of increased interest rates offered by a bank in the Philippines. Using the replication data, we replicate the results in Table 3 for the effect in the various treatment arms; the results for both the unconditional high interest rate and commitment "reward" interest rate treatment arms are statistically insignificant from 0. We then estimate the pooled treatment effect, using the variable for savings winsorized at 5% (since this is consistent with the winsorizing we perform in this paper). We divide by average weekly income of the treated (estimated using the replication data) times the 52-week duration of the study.

Kast, Meier, and Pomeranz (2016). This study looks at the effects of participating in a selfhelp peer group savings program in Chile. We use the intent-to-treat estimate of self-help peer groups on average monthly balance, 1817 pesos (Table 3, column 7). Although we would prefer to use the effect on ending balance, Figure 3b shows that average monthly balance is similar to ending balance. We use the estimate winsorized at 5% (since this is consistent with the winsorizing we perform in this paper). We divide the savings effect by average number of household members times average per capita household income in the treatment group (Table 1) times the 12-month duration of the study.

Kast and Pomeranz (2014). This study looks at the effects of removing barriers to opening savings accounts for low-income members of a Chilean microfinance institution, with a focus on the impacts on debt. Because of the focus on debt, we estimate the effect of treatment on *net* savings, or savings minus debt. To obtain estimates of the intent-to-treat effect, we multiply the average savings balance of active account users, 18,456 pesos, by the proportion of the treatment group who are active users (39%) and add the minimum balance of 1000 pesos times the proportion who take up but leave only the minimum balance (14%), all from Table 2. We then subtract the intent-to-treat effect on debt, -12,931 pesos. This gives an effect of $18,456 \cdot 0.39 + 1000 \cdot 0.14 - (-12,931) = 20,251.76$ pesos. We divide this by the average number of household members times average per capita monthly income (Table 1) times the 12-month duration of the study.

Prina (2015). This study looks at the effects of giving female household heads in Nepal access to savings accounts. We use the replication data to estimate the intent-to-treat effect on savings account balances after 55 weeks, the duration of the study. While the paper shows average bank savings among those who take up accounts, to estimate the intent-to-treat effect we take the bank savings variable and recode missing values (assigned to those who do not take up the account or are in the control group) as zero, then regress this variable on a treatment dummy. We divide by average weekly income among the treatment group from the endline survey (available in the replication data) times the 55-week duration of the study.

Sayinzoga, Bulte, and Lensink (2016). This study looks at the effects of offering financial literacy training to small farm owners in Rwanda. We use the replication data to estimate the intent-to-treat effect on the level of savings, replicating the point estimate in Table 2, panel (a) without covariates. We divide this by mean monthly expenditure for the treatment group multiplied by the

mean time between baseline and endline surveys, which we calculate directly in the replication data using survey date variables.

Schaner (2016). This study looks at the effects of offering very high, temporary interest rates in Kenya. We use the effect on bank savings (Table 3, column 2) and divide it by average monthly income of the treatment group (Table 4, column 6) times the 36-month duration of the study.

Seshan and Yang (2014). This study looks at the effects of inviting migrants from India working in Qatar to a motivational workshop that sought to promote better financial habits and joint decision-making with their spouses in India. The intent-to-treat effect on the level of savings comes from Table 3, column 1. We divide this by total monthly household income (constructed by adding the migrant's income and wife's household's income from Table 1, column 3) times the midpoint of the duration range of the study (15 months, since the study duration ranges from 13 and 17 months).

Somville and Vandewalle (2016). This study looks at the effects of defaulting payments into an account for rural workers in India. We use the effect of treatment on savings balances 23 weeks after the last payment, or 33 weeks after the beginning of the study (Table 5, column 3). We divide this by average weekly income (p. 20) times 33 weeks.

Suri and Jack (2016). This study looks at the effects of mobile money access in Kenya. The authors find that an increase in the penetration of mobile money agents within 1 kilometer of a household increases their log savings by 0.021 per agent for male-headed households and 0.032 per agent for female-headed households (Table 1). Using average savings of 286,752 shillings and the average change in agent density between 2008 and 2010 of 4.68 agents (Table S1), we thus calculate the effect of the average increase in access to mobile money on the level of savings as $(\exp(0.021) - 1) \cdot 286,752 \cdot 4.68$ for male-headed households, and $(\exp(0.032) - 1) \cdot 286,752 \cdot 4.68$ for female-headed households. We divide this by average daily per capita income (Table S1) times 365 days times 6 years times the average number of household members. We obtain the average number of household members from the replication data for Jack and Suri (2014) and estimate it separately for male- and female-headed households.