Banking on Trust: How Debit Cards Enable the Poor to Save More

Pierre Bachas
UC Berkeley
bachas@berkeley.edu

Paul Gertler
UC Berkeley
gertler@berkeley.edu

Sean Higgins
Tulane University
shiggins@tulane.edu

Enrique Seira
ITAM
enrique.seira@itam.mx

April 27, 2016

Abstract

Trust is an essential element of economic transactions, but trust in financial institutions is especially low among the poor, which may explain in part why the poor do not save formally. Debit cards provide not only easier access to savings (at any bank's ATM as opposed to the nearest bank branch), but also a mechanism to monitor bank account balances and thereby build trust in financial institutions. We study a natural experiment in which debit cards were rolled out to beneficiaries of a Mexican conditional cash transfer program, who were already receiving their transfers in savings accounts through a government bank. Using administrative data on transactions and balances in over 300,000 bank accounts over four years, we find that after receiving a debit card, the transfer recipients do not increase their savings for the first 6 months, but after this initial period, they begin saving and their marginal propensity to save increases over time. During this initial period, however, they use the card to check their balances frequently; the number of times they check their balances decreases over time as their reported trust in the bank increases. Using household survey panel data, we find the observed effect represents an increase in overall savings, rather than shifting savings; we also find that consumption of temptation goods (alcohol, tobacco, and sugar) falls, providing evidence that saving informally is difficult and the use of financial institutions to save helps solve self-control problems.

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 (Banfield, 1958; Knack and Keefer, 1997; Porta et al., 1997; Narayan and Pritchett, 1999; Algan and Cahuc, 2010). Trust is the "subjective probability with which an agent assesses that another . . . will perform a particular action" (Gambetta, 1988, p. 217). It is particularly important 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 et al., 2008), use of checks instead of cash (Guiso et al., 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 is low, however, as evidenced by the fact that majorities in 40 percent of countries included in the World Values Survey report lack of confidence in banks (Figure 1). Trust is especially low among the poor. In Mexico, for example, 71% of those with less than primary school report low trust in banks, compared to 55% of those who completed primary school and 46% of those who completed university (Figure 2). Along with fees and minimum balance requirements, trust is frequently listed as a primary reason for not saving in formal bank accounts (e.g., Dupas et al., 2016). At the country level, low trust in financial institutions is strongly correlated with the proportion of the population without bank accounts (Figure 3). Despite its importance, trust as a potential barrier to the poor saving in financial institutions has not been extensively studied (Karlan et al., 2014).[1]

Lack of trust in financial institutions may not be unfounded. Cohn et al. (2014) provide evidence that the banking industry fosters a culture of dishonesty relative to other industries. In Mexico in particular, bankers have been found to loot money by directing a large portion of bank lending

[1] Increased trust is proposed—but not explored further—as one channel through which no-fee savings accounts led to saving in Prina (2015).

to "related parties," i.e. shareholders of the bank and their firms (La Porta et al., 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 USD 200 million from savers (CNBV, 2014).[2] It is also telling that articles with financial advice in Mexican newspapers have titles like "How to Save for Your Graduation and Avoid Frauds" and "Retirement Savings Accounts, with Minimal Risk of Fraud." When contract enforcement is poor and fraud is rampant, trust becomes even more important (Guiso et al., 2004; Karlan et al., 2009) and people are understandably even more reluctant to use untrustworthy financial institutions (Bohnet et al., 2010).

While trust is important, it is not an innate characteristic but rather can be influenced by experience and information (Hirschman, 1984; Williamson, 1993; Attanasio et al., 2009). Debit cards (and mobile money) provide a low cost technology to monitor account balances and thereby build trust that a bank will neither explicitly steal deposits nor charge unexpectedly large hidden fees. Previous studies on debit cards and mobile money have focused on the effect of the lower transaction costs facilitated by these technologies to make purchases (Zinman, 2009), access savings and remittances (Suri et al., 2012; Schaner, forthcoming), and transfer money (Jack et al., 2013; Jack and Suri, 2014), but not their capacity to monitor and build trust in financial institutions. 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. Indeed, a lack of trust could explain why a number of randomized field experiments have found that even when take-up of accessible and affordable formal savings products is high, use is low in that most opened accounts have few transactions after the first 6 to 12 months (Ashraf et al., 2006; Dupas and Robinson, 2013a; Karlan and Zinman, 2014; Schaner, 2015).

We examine this hypothesis in the context of a natural experiment in which debit cards were rolled out geographically over time to beneficiaries of the Mexican conditional cash transfer program Oportunidades. The beneficiaries had been receiving their transfers into savings accounts for five

[2] This type of fraud is not uncommon: we scraped the online news archives of all electronic newspapers and news websites we could find in Mexico (129 total) using several keywords, then filtered the results by hand to keep only relevant stories. We found 1338 news stories associated with savings fraud in 2014 and 2015 alone.

years on average before debit cards were attached to their accounts, but typically did not use the accounts to save as they immediately withdraw most if not all of the transfer.[3] The phased geographic rollout provides plausibility exogenous variation in assignment of debit cards to beneficiaries in a difference in difference context. For the analysis, we use high frequency administrative data on bank transactions for over 340,000 beneficiary accounts in 370 bank branches over 4 years as well as several household surveys of a sample of the same beneficiaries.

Using the high frequency administrative data, we find that beneficiaries initially used debit cards to check account balances without any increase in savings, but over time the frequency of account balance checks fell and savings rates rose. We estimate that after one year, the share of total income saved each payment period increased by 5 percentage points and that after nearly two years those with cards saved 8 percentage points more per period.

The delayed initiation of savings suggests some kind of learning. We explore three kinds of learning that may be occurring: (i) learning to trust the bank, (ii) learning to use the debit cards and ATMs, and (iii) learning that the program will not drop beneficiaries who accumulate savings. Using household survey data, we find support for the "learning to trust" hypothesis but not for the other two types of learning. Specifically, we find that 27 percent of beneficiaries who have had the debit card for less than 6 months report that they do not trust the bank, compared to just 17 percent of those who have had the card for more than 6 months. We find very few beneficiaries who report not knowing how to use the technology or fear the program will drop them if they accumulate savings, and no change over time comparing those that have had the debit card less than and more than 6 months. We also find that those who have had the card more than 6 months report checking their balances significantly less frequently than those who have had the card less than 6 months, consistent with our finding from administrative data that when beneficiaries first get the debit card, they check their balances often, but the frequency of checking falls over time.

We then test whether the increase in the bank account balances is an increase in total savings or a substitution from other forms of saving, both formal and informal. Using panel household survey data, we find that after one year the treatment group increases total savings by about 5 percent of income relative to the control group, which is close in magnitude to the effect we see in the

[3]This is consistent with findings from other countries such as Brazil, Colombia, and South Africa, in which cash transfers are paid through bank accounts, but recipients generally withdraw the entire transfer amount each pay period and do not save in the account (Bold et al., 2012).

administrative account data. We find no differential change in income or assets in the treatment group compared to the control. These results suggest that the increase in saving is not driven by higher income but by (voluntarily) lowering current consumption and that that the increase in bank savings does not crowd out other forms of saving (consistent with Ashraf et al., 2015; Dupas and Robinson, 2013a; Kast et al., 2012).

Finally, a portion of the increase in savings in achieved through a decrease in the consumption of alcohol, tobacco, and sugar—the most frequently mentioned temptation goods in Banerjee and Mullainathan (2010). Indeed, this is the only consumption category with a statistically significant decrease after receiving the card. Although the poor do save via cash at home (Collins et al., 2009), saving informally is harder as the "money is hot" and susceptible to temptation spending, either by the beneficiary herself or by her husband if she lacks control over his access to her savings (Ashraf, 2009). Indeed, we also find that among beneficiaries living with a spouse or partner, those with lower baseline bargaining power relative to their spouse have a higher increase in savings after receiving the debit card. Our results suggest that saving in formal financial institutions may help solve some of the intra-household bargaining and self-control problems associated with trying to save informally.

These results are important for public policy as building savings in formal financial institutions has positive welfare effects for the poor and nearly half of the world's adults do not use financial institutions (Demirgüç-Kunt et al., 2015). The poor have used savings products to decrease income volatility (Chamon et al., 2013), accumulate money for microenterprise investments (Dupas and Robinson, 2013a), invest in preventative health products and pay for unexpected health emergencies (Dupas and Robinson, 2013b), and invest in children's education (Prina, 2015). Various randomized experiments have found that providing affordable and accessible savings accounts to the poor increases their future agricultural/business output and household consumption (Brune et al., 2016; Dupas and Robinson, 2013a), decreases debt (Kast et al., 2012; Atkinson et al., 2013), and improves their ability to cope with shocks (Prina, 2015). For these reasons, Mullainathan and Shafir (2009) conclude that access to formal savings services "may provide an important pathway out of poverty."

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 are already embarking on digitizing their cash transfer payments through banks and mobile money. Furthermore, the technologies of debit cards and ATMs or point of sale (POS) terminals—which can be used to check balances and access savings—are simple, prevalent, and potentially scalable to millions of government cash transfer recipients worldwide.

2 Institutional Context

We examine the the rollout of debit cards to urban beneficiaries of Mexico's conditional cash transfer program Oportunidades whose benefits were already being deposited directly into 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 (e.g., Gertler, 2004; Parker and Teruel, 2005). The program provides bimonthly cash transfers to poor families in Mexico, seeking to alleviate poverty in the short term and break the intergenerational poverty cycle in the long term by requiring families to invest in the human capital of childen by sending their children to school and having health check-ups. 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).

Oportunidades opened savings accounts in banks for a portion of beneficiaries in 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 lowers both the ability of corrupt local officials to skim off benefits 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, (iii) decrease robberies and assaults of program officers and recipients transporting cash on known days, and (iv) increase the financial inclusion of poor households. By the end of 2004, over one million families received their benefits directly deposited into savings accounts in Bansefi, a government bank created to increase savings and financial inclusion of underserved populations (Figure 4).[4]

[4]Originally 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.

The Bansefi savings accounts have no minimum balance requirement or monthly fees and pay essentially no interest.[5] 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, many beneficiaries live far from their nearest branch, meaning that accessing their accounts involved large transaction costs for many beneficiaries. Overall, the savings accounts were barely used prior to the introduction of debit cards. In 2008, the year before the rollout of debit cards, the average number of deposits per bimester[6] was 1.05 including the deposit from Oportunidades, the average number of withdrawals was 1.02, and 98.9 percent of the transfer was taken during the first withdrawal following payment.

In 2009, the government announced that they would issue Visa debit cards to beneficiaries that were receiving their benefits directly deposited into Bansefi savings accounts. The cards enabled account holders to withdraw cash from, make deposits into, and check balances of their account at any bank's ATM as well as make electronic payments at any store accepting Visa. The cards included two free ATM withdrawals every bimester at any bank's ATM, after which ATM withdrawal fees averaged 13 pesos (about $1 using 2009 exchange rates) but varied by bank.

Opportunities used direct deposit into savings accounts for its beneficiaries in 275 out of Mexico's 550 urban localities. Of these, debit cards were rolled out to approximately 100,000 beneficiaries in 143 localities in 2009 (wave 1) and to an additional 75,000 beneficiaries in 88 localities in late 2010 (wave 2). Another 170,000 beneficiaries in the remaining localities were scheduled to receive cards between November 2011 and February 2012 (control group) after the end date of our data period. The map in Figure 5 shows that the treatment and control waves had substantial geographical breadth and that some treatment and control localities were physically close.

The sequence with which localities switched was determined as a function of the proportion of households in the locality that were eligible for the program but were 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 1 compares the means of locality-level variables and account-level variables from the control, wave 1, and wave 2 localities using data from the population census from 2005, poverty estimates from Oportunidades from 2005, Bansefi branch locations from 2008,

[5]Nominal Interest rates were between 0.09 and 0.16 percent per year compared to an inflation of around 5 percent per year during our sample period.

[6]The 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.)

and 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 percent of the transfer withdrawn also differs (it is lower in wave 1 than the control and insignificantly different but with a higher point estimate in wave 2), but is high in all cases (ranging from 97.5 percent to 99.6 percent of the transfer), indicating very low savings in the account prior to receiving the card. In Sections 4 and 9.1, we will test and show that trends of saving, income and consumption were parallel across waves.

3 Data

We use a rich combination of administrative and survey data sources. To examine the effect of rollout of the debit cards on savings we use administrative data from Bansefi at the account level for 342,709 accounts at 380 Bansefi branches for a four-year period, from November 2007 to October 2011. These data include the bimonthly transfer amount the timing and amount of transactions made in the account, bimonthly average savings balances, the date the savings account was opened, and the month the card was awarded to the account holder. The average account had been opened 5.3 years before getting the card.

To test whether the delayed savings effect and increasing propensity to save over time can be explained by learning to use the technology, learning the program rules, or building trust in the bank, we use the Survey of Urban Households' Sociodemographic Characteristics (ENCASDU), conducted by Oportunidades at the end of 2010. We also use the Payment Method Survey, a household survey conducted by Oportunidades in 2012 aimed at eliciting satisfaction with and use of the debit cards.

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 early 2010. This survey has comprehensive modules on consumption, income, and assets. We merge these data with administrative data from Oportunidades on the transfer histories for this sample—which we use to add transfer income into total income and to identify which households are Oportunidades recipients, given the common misreporting of transfer receipt in surveys (Meyer et al., 2015)—and on the dates that debit cards were distributed in each locality.

Because the final pre-treatment wave of ENCELURB in 2004 is five years prior to wave 1 of the debit card rollout, we supplement our parallel trends test in ENCELURB with data for the intervening period (2004-2008) from the National Household Income and Expenditure Survey (ENIGH), a repeated cross-section; we merge the publicly available ENIGH with restricted-access locality identifiers provided by the National Institute of Statistics and Geography (INEGI) to determine which surveyed households were in treatment and control localities, and restrict the analysis to the poorest 20 percent of surveyed households to proxy for Oportunidades recipients.

Figure 4a shows the timing of the administrative Bansefi account balance and transaction data, while Figure 4b shows the timing of the household survey data (merged with additional administrative data) we use, both relative to the rollout of debit cards.

4 Effect of Debit Cards on Stock of Savings

Figure 6 presents average balances over time; even the raw data are very telling. Panel (a) compares the first wave of debit card recipients to the control group, with a dashed vertical line indicating the time when wave 1 localities received debit cards, while Panel (b) compares the second wave to the control, with a dashed vertical line indicating the time when wave 2 localities received debit cards. Strikingly, average balances increase sharply for the first wave after receiving the card, but the effect is not immediate: it begins about four bimesters after receiving the card and the larger increase happens after a year with the card. By October of 2011, wave 1 has average balances of around 2000 pesos, over three times that of the control group. Average balances also increase over time with the card in wave 2, although we have information for less bimesters after wave 2's later switch to debit cards.

Although our data on average balances is by bimester, some payments get shifted to the end of the prior bimester, so we group adjacent bimesters into four-month periods for the remainder of the analysis. Because we have four years of data, this leaves us with 12 four-month periods. To compare the stock of savings in the treatment and control groups while controlling for individual observables and unobservables, as well as any common time shocks, we use a period-by-period difference-in-differences (DID) strategy and estimate:

$$Balance_{it} = \lambda_i + \delta_t + \sum_{k=1}^{12} \phi_k T_{j(i)} \times \mathbb{I}(t = k) + \varepsilon_{it} \tag{1}$$

where $Balance_{it}$ is the average balance in account i over period t (specifically, end of day balances were averaged over the number of days in the bimester by Bansefi, and we average the average balances over the two adjacent bimesters that make up the four-month period). Following other papers measuring savings (e.g., Kast et al., 2012; de Mel et al., 2013; Karlan and Zinman, 2014; Chetty et al., 2014; Akbas et al., 2015; Dupas et al., 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.[7] The λ_i are account-level fixed effects which control for observable and unobservable time-invariant characteristics of the beneficiaries, δ_t are time-period dummies that control for general macro trends such as bimester-specific shocks that affect both treatment and control groups, $T_{j(i)} = 1$ if locality j in which account holder i lives is a treatment (i.e., wave 1 or 2) locality, and $\mathbb{I}(t = k)$ are time period dummies. 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 estimate cluster-robust standard errors, ε_{it}, clustering by Bansefi branch. Since one time period dummy must be omitted from (1), we follow the standard procedure of omitting the four-month period immediately preceding the change to cards. We estimate (1) separately for wave 1 and wave 2.

The coefficients of interests are the ϕ_ks, which measure the average difference in balances between the control and treatment group in bimester k. The raw data clearly suggest that pre-treatment trends of savings were parallel across control and treatment groups before getting the card; we test this statistically by testing $\phi_1 = \cdots = \phi_{\ell-1} = 0$ where ℓ is the period of switch. (In wave 1, ℓ is the period November 2009-February 2010, and in wave 2 it is the period November 2010-February 2011.) Figure 7 plots the ϕ_ks and shows that pretreatment coefficients are, in most periods, not individually different from zero, and we cannot reject that pre-trends are equal between treatment and control: the p-value for the F-test of $\phi_1 = \cdots = \phi_{\ell-1} = 0$ is 0.823 for wave 1 and 0.110 for wave 2.

The cards also led to an increase in use of the accounts: even though clients' average number

[7] To avoid truncating a true treatment effect or time trend, we winsorize within each wave and time period.

of deposits remained constant at one, corresponding to the Opportunidades transfer, their number of withdrawals increased. Figure 8 plots the number of withdrawals per Oportunidades deposit per bimester: prior to receiving the debit card, both the treatment and control groups received a single deposit and made a single withdrawal, and hence the number of withdrawals per deposit was stable at very close to 1. After receiving the card, clients made on average 1.4 withdrawals per deposit, per bimester.[8] Figure 9 shows the distribution of the number of withdrawals, before and after receiving the card. Prior to receiving the card, over 90% of clients just made a single withdrawal. 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. This immediate increase in use of the account after a decrease in the transaction costs of accessing money agrees with the prediction of the Baumol (1952) and Tobin (1956) model of money demand in the face of transaction costs, and with empirical evidence that ATMs and debit cards lead to reduced transaction costs and an increased number of withdrawals (Attanasio et al., 2002; Alvarez and Lippi, 2009).

This increased account use will also lead to a "mechanical" increase in our dependent variable, average balance, because beneficiaries will be leaving a portion of their transfer in the account for a longer period of time. For example, the 27% who make two withdrawals with the card withdraw during the first withdrawal 71% of the total amount withdrawn over the bimester (which might be less than the total deposited if they intend to save some), then return on average 9 days later to make a second withdrawal of the remaining 29% of the total they withdraw over the period. For these 9 days, 29% of the amount they withdrew over the bimester (and hence did not save) is nevertheless captured in the balance; we call the effect of this on the average balance over the period the "mechanical effect." Furthermore, even for those who make one withdrawal of the entire transfer, the average balance will be positive if they wait some number of days after receiving the deposit before withdrawing it.[9] We compute the mechanical effect for each account in each bimester

[8] The reason we measure withdrawals per deposit, rather than simply withdrawals, is that some beneficiaries receive two payments in some bimesters.

[9] From discussions with Bansefi and Oportunidades officials, a portion of the lapse of time that the transfer remains in the account before being withdrawn can be explained by an administrative feature of the program: Opportunidades provides transfer recipients with a calendar of dates when their accounts will be credited; however, Opportunidades often transfers the deposits to the Bansefi accounts several days prior to the due date for two reasons. First, it does so in order to avoid backlogs and ensure that all accounts have been credited by the announced calendar dates. Second, it transfers funds in bulk (all on the same day for a particular locality) but staggers the calendar dates on which Oportunidades beneficiaries are told their funds will be made available, in order to avoid congestion at ATMs and

using data on the amounts of and timing between deposits and withdrawals during each bimester, as described in detail in Appendix B, and subtract this from the average balance to create a variable we call "net balance."

Figure 10 shows the ϕ_k coefficients from (1), using net balance (i.e., average balance minus mechanical effect) as the dependent variable. The debit cards lead to an increase in the stock of savings, with net balances in the account tending to increase over time with the debit card. In Wave 1, there is a marked delay of about one year before beneficiaries start using the account to save. As expected, after subtracting out the mechanical effect from average balances, the treatment effect is smaller in magnitude, reaching about 900 pesos after two years with the card, compared to 1400 pesos in the average balances specification.

5 Effect of Debit Cards on Marginal Propensity to Save

To measure the propensity to save, we control for the amount received in transfers each period. This is important since there is a large amount of variation in transfers received within accounts over time, as well as between accounts. The variation within an account over time can be explained by local elections in certain localities, compliance with program conditions, payment amounts varying depending on the time of year, and family structure.[10]

In the spirit of asset accumulation models, we assume that savings in period t is a function of assets in period $t-1$, income in period t, (time-invariant) individual preferences, and period-specific shocks such as changes to prices:

$$Savings_{it} = f(Assets_{i,t-1}, Income_{it}, \lambda_i, \delta_t). \tag{2}$$

Linearizing f, separating assets into the savings balance in the Bansefi account and other assets,

Bansefi branches.

[10]When there is an election, Oportunidades has 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 toward the end of 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. As an example of payments varying by time of year, 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. Changes in family structure affect the transfer amount because one child might age into or out of the program, for example.

and separating income into Oportunidades transfer income and other income gives

$$\begin{aligned} Savings_{it} = \lambda_i + \delta_t + \beta Net\ Balance_{i,t-1} + \kappa Other\ Assets_{i,t-1} \\ + \gamma Transfers_{it} + \xi Other\ Income_{it} + \varepsilon_{it}, \end{aligned} \tag{3}$$

where ε_{it} captures period-specific idiosyncratic shocks. Our adminitrative data from Bansefi only include transfers and balances, but not other income and other assets; after removing these terms from (3), each household's average other income and average assets over time are captured by the fixed effect λ_i, while idiosyncratic changes in these variables over time add noise in the error term. Our measure of savings at time t is the difference in net balance between time t and time $t-1$; we thus have

$$Net\ Balance_{it} - Net\ Balance_{i,t-1} = \lambda_i + \delta_t + \beta Net\ Balance_{i,t-1} + \gamma Transfers_{it} + \varepsilon_{it}, \tag{4}$$

where γ gives the marginal propensity to save out of transfer income. Since transfers are on average about 20% of total income in our sample, dividing our estimates by five gives a rough approximation of the marginal propensity to save out of total income. Grouping terms in (4) gives

$$Net\ Balance_{it} = \lambda_i + \delta_t + \theta Net\ Balance_{i,t-1} + \gamma Transfers_{it} + \varepsilon_{it}, \tag{5}$$

where $\theta = 1 + \beta$; then to estimate the effect of receiving a debit card on the marginal propensity to save out of transfers and allow this effect to change over time, we estimate

$$\begin{aligned} Net\ Balance_{it} = \lambda_i + \delta_t + \theta Net\ Balance_{i,t-1} + \sum_{k=2}^{12} \alpha_k T_{j(i)} \times \mathbb{I}(k=t) \\ + \sum_{k=2}^{12} \gamma_k Transfers_{it} \times \mathbb{I}(k=t) + \sum_{k=2}^{12} \psi_k Transfers_{it} \times T_{j(i)} \times \mathbb{I}(k=t) + \varepsilon_{it}, \end{aligned} \tag{6}$$

winsorizing both net balance and transfers at the 95th percentile.

As is well-known, however, fixed effects panel data models with a lagged dependent variable (also known as dynamic panel data models) are biased and inconsistent (Nickell, 1981). We thus use the system generalized method of moments (GMM) estimator proposed by Arellano and Bover

(1995) and Blundell and Bond (1998), which is consistent for fixed T, large N (as we have here) and performs well in Monte Carlo simulations, especially for large N (Blundell et al., 2001). The two-step system GMM estimator also appears to perform better than Kiviet's (1995, 1999) and Bruno's (2005) least squares dependent variable correction methods when N is large (Bun and Kiviet, 2006). The effect of the debit card on the marginal propensity to save out of transfer income in bimester k is $\alpha_k/\mu_k + \psi_k$, where μ_k is average transfers in bimester k; Figure 11 plots the $\alpha_k/\mu_k + \psi_k$ estimates along with their confidence intervals. Standard errors of the parameters in (6) are clustered at the bank branch level and corrected for finite sample bias following Windmeijer (2005); the formula for the variance of $\alpha_k/\mu_k + \psi_k$ is then approximated using the delta method. As before, we estimate (6) separately for wave 1 and wave 2.[11]

In Figure 11, the marginal propensity to save out of the transfer is not significantly different between the treatment and control prior to receiving the card, and we observe a delayed effect after receiving the card: in wave 1, the effect remains statistically insignificant from 0 for the first three 4-month periods after receiving the card, while in wave 2 it is insignificant from 0 for the first two periods after they receive the card. The MPS then increases over time and, in wave 1 where we have more post-treatment data, increases substantially over the two years with the card. After one year with the card (in the November 2010–February 2011 period), account-holders save 26.8% of their transfer, which—using household survey data merged with administrative data from Oportunidades on bimonthly transfers to determine the proportion of total income coming from transfers—equals about 5.4% of total income. After close to two years (in the July–October 2011 period), it equals 39.1% of the transfer, or 7.9% of total income. In wave 2, the MPS increases sooner, reaching 10.8% of the transfer or 2.2% of total income after between 6 months and one year with the card.

[11]Following the best reporting practices outlined in Roodman (2009a), 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; Roodman, 2009b), to reduce the number of instruments we only use one lag of $Balance_{i,t-1}$ as an instrument. Because $Transfers_{it}$ is predetermined but not strictly exogenous, variables on the right hand side of (6) 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 (6) 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 70. 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.

6 Mechanisms

Why do we see a delayed savings effect after receiving the debit card, and why does the marginal propensity to save out of the transfer gradually increase with time? We conjecture that learning is at play and explore three kinds of learning: operational learning (i.e., learning how to use the technology), learning the program rules (specifically, that the program will not drop beneficiaries who accumulate savings), and learning to trust (that the bank is a safe place to save). The first involves knowledge of how to use the debit card and ATM, memorizing the card's PIN, etc. The second involves learning that the program will not use accumulated savings as a signal that the family is actually not poor enough to be receiving Oportunidades benefits. These first two explanations were conjectured by Oportunidades program officials when we shared our initial results from the administrative Bansefi data. The third involves learning that the risk of getting the money "stolen" in the form of hidden fees, operational errors, or nefarious behavior by the bank is lower than initially believed. We find evidence that beneficiaries use the card to check their account balances, and that it thus provides them with a technology to monitor bank behavior, ensure that their money is not disappearing, and subsequently build trust in the bank.

We first use data from the ENCASDU, a survey that directly asks beneficiaries "Do you leave part of the monetary support from Oportunidades in your bank account?" and, if the response is no: "Why don't you keep part of the monetary support from Oportunidades in your Bansefi bank account?" The second question includes pre-written responses and an open-ended response. An example of an answer coded as lack of knowledge is "They didn't explain the process for saving." An example of an answer coded as fear of being dropped from the program is "Because if I save in that account, they can drop me from the program." An example of an answer coded as lack of trust is "Because if I don't take out all the money I can lose what remains in the bank." The ENCASDU surveyed 8788 Oportunidades beneficiary households across rural, semi-urban, and urban areas; of these, the 1674 that received Oportunidades benefits in savings accounts tied to debit cards at the time of the survey make up our sample.

We estimate

$$y_i = \alpha + \gamma \mathbb{I}(\text{Card} \leq 6 \text{ months})_i + u_i, \tag{7}$$

where three regressions are run in which the dependent variable $y_i = 1$ if the beneficiary reports not saving due to (i) a lack of knowledge, (ii) fear they will be dropped from the program, or (iii) lack of trust. We estimate the unconditional probability, i.e. beneficiaries who report saving are included in the regression with $y_i = 0$. The unconditional probability is the more relevant measure; instead using the conditional probability (only including those who save in the regression) would mean that the delayed effect we have observed of debit cards on savings could drive the result. Standard errors are clustered at the locality level. We test the null hypothesis $\gamma = 0$, where a rejection of the null would imply that the dependent variable we are testing—which is related to either learning to use the technology, learning program rules, or learning to trust the bank—changes over time with the card. Although this survey is cross-sectional, we exploit the variation in time with the debit card, exogenously determined by the staggered locality-level rollout of the cards.

Figure 12a and Table 2a show the results. The first thing to note is that lack of knowledge and fear of being dropped from the program after saving are rarely cited as reasons for not saving (combined, less than 4 percent of the sample who have had the card for less than 6 months do not save for these reasons), while lack of trust is cited by 27 percent of those who have had the card for less than 6 months. Second, the proportion who report not saving due to a lack of knowledge does not change over time; in contrast, trust increases gradually with experience: beneficiaries with more than 6 months with the card are 36 percent less likely to report not saving due to low trust than those with less than 6 months with the card.

Next, we explore mechanisms behind operational learning and learning to trust the bank using the 2012 Payment Methods Survey. The survey includes a number of questions related to operational learning: "What have been the main problems you have had with the ATM?"; "In general, does someone help you use the ATM?"; "Do you know your PIN by heart?"; "Did they tell you that with the card you have a Bansefi savings account?" It also includes a question on balance checking ("In the last bimester, how many times did you check your balance?"), which is a mechanism that beneficiaries could use to build trust in the bank once they have a debit card. The Payment Methods Survey included 5381 households, drawn by stratified (by payment method and locality) random sampling from all Oportunidades beneficiaries; of these, our sample is made up of the 1641 who received their benefits on debit cards tied to savings accounts.

We again use specification (7), with y_i equal to: (i) the self-reported number of balance checks over the past bimester; (ii) the self-reported number of balance checks over the past bimester without withdrawing any money, constructed as the total number of balance checks minus the number of withdrawals; and dummies if the respondent reports (iii) it is hard to use the ATM; (iv) she gets help using the ATM; (v) she knows her PIN; (vi) she knows she can save in the account. Because this survey was conducted in 2012, those with the card for at least 6 months now include both wave 1 and wave 2, while beneficiaries in the localities we treat as control localities throughout this paper make up the group with cards for less than 6 months.

Figure 12b and Table 2b show the results. Both the number of balance checks and number of balance checks without withdrawing decrease over time with the card. Making trips to the ATM specifically to check the account balance (i.e., making a balance check without withdrawing any money) decreases by 36 percent after six months compared to the first 6 months (from an average of 0.53 balance checks without withdrawing to an average of 0.34), while most measures that indicate knowledge of how to use the technology do not change over time: the proportions who report it is hard to use the ATM (around 10 percent), that they get help using the ATM (55 percent), and that they know they can save in the account (32 percent) do not change, although there is a statistically significant increase in the proportion who know their PINs (from 49 to 58 percent).

Finally, we use the administrative transactions data from the 342,709 Bansefi accounts, which include the date, time, and fee charged for each balance check at an ATM for each account, to investigate whether the mechanism that appears to be driving the increase in trust—balance checks which clients use to monitor and, over time, build trust in the bank—holds true in the administrative data; the increased power we have from a large number of observations in the administrative data allows us to take a more granular look at balance checks over time. Note that balance checks at a Bansefi branch are possible both before and after receiving the debit card. Nevertheless, if the distance to the nearest bank branch is high, the debit cards provide a technology that greatly reduces the cost of balance checking (by enabling clients to check their balances at the closest ATM of any bank). Since balance checks at a Bansefi branch are not charged a fee—unlike balance checks at ATMs in Mexico—we do not observe them in our data, which is why average balance checks (at ATMs) in the graph begin after debit card receipt.

Figure 13 plots the number of times people check their balance per bimester, with vertical lines indicating the timing of card receipt. Again due to the shifting of some payments to the end of the previous bimester (which might affect the bimester timing of balance checks), we continue grouping adjacent bimesters into four-month periods. We observe that the number of balance checks per bimester is initially around 2.5 checks on average in wave 1 and 1.5 checks on average in wave 2, but in both waves decreases during each four-month period after beneficiaries receive the card, consistent with the trust-building hypothesis. This is consistent with the learning to trust hypothesis: through checking her balance, the client learns that her money is still in the account and updates downward her belief about the risk of losing money. With simple Bayesian learning, balance checking has decreasing marginal benefit and she therefore checks her balance less over time.[12]

7 Increase in Overall Savings vs. Substitution

The increase in formal Bansefi account savings might come at the expense of other types of savings that the household is already conducting, in such a way that total savings is not affected. The question of whether the observed increase in Bansefi savings crowds out other saving is relevant not only if one is concerned with total household savings, but also to understand the mechanics through which the effect on formal savings is operating and as a first step towards thinking about the broader welfare implications of providing a formal savings account with a debit card.

Does the provision of the debit card and the resulting increase in formal savings represent an increase in *overall* savings, or is it merely a substitution from other forms of saving? To address this question, we use Oportunidades' ENCELURB panel survey, conducted 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 in urban and semi-urban areas.[13] Of the 6272 households in the post-treatment wave of ENCELURB, 2951 live in urban areas and, according to administrative data provided by Oportunidades and merged with the survey, are Oportunidades beneficiaries when interviewed in the post-treatment wave and receive their benefits in a savings account (with or without a debit card); this is the sample used in our analysis, except in the placebo tests described in Section 9.

[12]The operational learning hypothesis makes a different prediction regarding the evolution of balance checks over time: since it becomes easier—less costly—for a beneficiary to check her balance as she learns to use the technology (e.g., by memorizing her PIN or learning how to use the ATM), if anything we might expect her to check her balance more over time if operational learning were the mechanism at play.

[13]The 2002, 2003, and 2004 waves had around 17,000 households, but due to budget constraints the number of localities was cut for the 2009-2010 wave. The consumption, income, and assets modules of Oportunidades' analogous survey for rural areas have been used by Angelucci and De Giorgi (2009), Attanasio et al. (2013), de Janvry et al. (2015), Gertler et al. (2012), and Hoddinott and Skoufias (2004), while these modules from the ENCELURB have been used by Angelucci and Attanasio (2013) and Behrman et al. (2012).

As before, we use a DID strategy where we examine changes in consumption, savings, and income 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 have not yet received cards (waves 2 and control), respectively referring to them as "treatment" and "control" in this section of the paper. The identification assumption is that in the absence of the debit card, treatment and control groups would have experienced similar changes in consumption, income, and assets. We formally test for parallel trends in Section 9, and since we indeed find that trends were parallel prior to treatment, we now test whether there was an increase in savings, which we construct as income minus consumption from the income and consumption modules of ENCELURB. We estimate

$$y_{it} = \lambda_i + \delta_t + \gamma D_{j(i)t} + \nu_{it} \tag{8}$$

separately for five dependent variables: consumption, income, savings (constructed as income minus consumption), purchase of durables, and an asset index. All variables except assets are measured in pesos per month, i indexes households, and t indexes survey years. As before, all variables are winsorized at the 5% level (i.e., the 95th percentile, and also the 5th percentile in the case of variables that do not have a lower bound of 0) to avoid results driven by outliers. The 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. 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; i.e., in the notation used in specifications (1) and (6), $D_{j(i)t} \equiv T_{j(i)} \times \mathbb{I}(t = 2009\text{–}10)$. We use the locality of residence prior to treatment to avoid capturing migration effects in our estimation and estimate cluster-robust standard errors at the locality level.

If the increase in formal savings constitutes an increase in total savings then we expect $\gamma > 0$ for total savings (defined as income minus consumption), and if we observe $\gamma = 0$ for income we expect $\gamma < 0$ for consumption. 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). This is indeed what we find. Figure 15 shows that consumption decreased by about 130 pesos on average (statistically significant at the 10% level). Meanwhile, there is no effect on income; we 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 ($p = 0.072$). Purchase of durables and the stock of assets do not change, ruling out a crowding out of these forms of saving. The increase in savings, measured as income minus consumption—which, although a crude measure of savings, is commonly used (e.g., Dynan et al., 2004)—is estimated at slightly more than 200 pesos, and is significant at the 5% level.

The results in Figure 15 are from our main specification where we winsorize the dependent variable at 5 percent (specifically, at the 95th percentile, as well as the 5th percentile if the variable does not have a lower bound of 0). Appendix Table A1, columns 1–3 show that the effects are robust to using the raw data without winsorizing, winsorizing at 1 percent, or—as in our main specification—winsorizing at 5 percent (we follow Kast and Pomeranz (2014) and others who show the robustness of results to these three possibilities). They are also robust to including baseline characteristics interacted with time fixed effects, as well as municipality-specific time effects, both to control for specific time trends more flexibly (Appendix Table A1 columns 4 and 5).[14]

These results mean that total savings—not just formal savings—increase, and that at least a substantial portion of this increase in being funded by lower consumption. A back of the envelope calculation reveals that the magnitude of the increase in monthly savings from this household survey is in line with the average increase of savings in the account from the administrative data: from the propensity to save specification, after 1 year, beneficiaries who received cards in wave 1 save 26.8% of their transfer more than the control group. Using ENCELURB, transfers are, on average, 20.2% of income for the treatment group, implying that the savings effect in the Bansefi administrative data

[14] The household characteristics interacted with time fixed effects in this robustness check are measured at baseline and include characteristics of the household head (whether the household head worked, a quadratic polynomial in years of schooling, and a quadratic polynomial in age), whether the household has a bank account, variables used to measure poverty by Oportunidades (the proportion of household members with health insurance, the proportion aged 15 or older that are illiterate, the proportion aged 6 to 14 that do not atted school, the proportion aged 15 or older with incomplete primary education, and the proportion aged 15 to 29 with less than 9 years of schooling), and dwelling characteristics (dirt floor, no bathroom, no water, no sewage, number of occupants per room).

is about 5.4% of income. The effect for savings (income minus consumption) in the ENCELURB household survey data shown in Figure 13 equates to 4.9% of income. Taken at face value, this suggests that most of the increase in savings in the account is new saving. This result is consistent with other studies where formal savings products were offered, which found that the increased savings in these products did not crowd out other forms of saving (Ashraf et al., 2015; Dupas and Robinson, 2013a; Kast et al., 2012).

8 Does Money Burn a Hole in Your Pocket?

Because the accounts pay no interest, but there was clearly an unmet demand for savings among program beneficiaries, we explore why they were not able to save before (for example, under the mattress). Since the results in Figure 15 show that the debit card induces higher total savings through decreased consumption, we might expect that it influences different components of consumption differentially. We thus examine the proportion of income spent on several categories of consumption goods, listed in order of average budget share: meat, dairy, and produce; tortillas and cereals; health and education; transportation; fats and sweets (junk food, fats, soda); temptation goods (where we group the three most frequently cited temptation goods in Banerjee and Mullainathan (2010): alcohol, tobacco, and sugar); and entertainment. We use the proportion of total income spent on each consumption category, rather then the level of consumption in that category, because individual × time-specific shocks to income, which we expect to be passed through as shocks to various consumption categories, would otherwise add noise to the estimation through the error term; we use total income rather than total consumption in the denominator because, from the results in Figure 15, total income does not change differentially between the treatment and control groups.

We estimate a DID specification with household and year fixed effects and standard errors clustered at the locality level; specifically, for each consumption category g,

$$\frac{Consumption_{git}}{Income_{it}} = \lambda_{gi} + \delta_{gt} + \gamma_g D_{j(i)t} + \nu_{git}, \tag{9}$$

where $Consumption_{git}$ is monthly consumption of goods category g by household i at time t (in pesos) and $Income_{it}$ is total monthly income of household i at time t (in pesos). We find that the only category of goods in which the treatment group shows a statistically significant decrease relative to the control is that of temptation goods (Figure 16).[15] More specifically, the treatment group reduces the proportion of income it spends on temptation goods by 20% relative to the control group. Nevertheless, as the thick horizontal bars in Figure 16 show, and as we would expect, only a small portion of income is spent on temptation goods at baseline (3%). As a result, the statistically significant reduction in spending on temptation goods only explains 16% of the total effect of the debit cards on consumption. We refrain from attempting to determine what categories explain the remaining 84% of the effect since the results for other spending categories are not statistically significant.

Although our grouping of temptation goods is based on the goods most frequently mentioned by Banerjee and Mullainathan (2010), it could be viewed as arbitrary (and, indeed, we do not find a decrease in the grouping of fats and sweets—junk food, fats, and soda—which could also be classified as temptation goods); we thus look separately at each item in the temptation good category, and find a statistically significant decrease in consumption of alcohol and sugar, but not of tobacco.

We interpret this result as evidence that it is difficult to save informally due to self-control problems, and that these problems can be partially solved by access to a formal savings account (but that low indirect transaction costs and trust in the bank are necessary conditions for these formal savings accounts to be used). This finding is consistent with the demand for commitment savings devices (e.g., Ashraf et al., 2006; Bryan et al., 2010) if the savings accounts without debit cards, which could be used as an even stronger commitment due to the high indirect cost of accessing savings, would have been too strong of a commitment (since strong commitment devices have low take-up and low use relative to weak commitment devices: see Karlan and Linden, 2014; Laibson, 2015), or if the bank accounts were merely not trusted prior to being able to cheaply monitor them with debit cards. Under either explanation, trust appears to have been a necessary condition for formal saving, given the delayed savings effect and self-reported reasons for not saving initially.

The self-control problems that prevent the poor from saving prior to having access to a trusted formal savings account could result directly from an asset-based poverty trap, as in Bernheim et al.

[15] The whiskers in Figure 16 show the 95% confidence intervals when no adjustment is made for multiple hypothesis testing. After adjusting for multiple hypothesis testing using the sharpened false discovery rate (Benjamini et al., 2006; Anderson, 2008), the result for temptation goods is significant at the 10% rather than 5% level ($p = 0.023, q = 0.086$).

(2015), a model that is consistent with the empirical finding that microcredit decreases temptation good consumption (Angelucci et al., 2015; Augsburg et al., 2015; Banerjee et al., 2015). Alternatively, it is possible that the self-control problems stem from the timing of access to the money: Carvalho et al. (2016), using exogenous variation in the timing of an experiment relative to payday, find that those who are more financially constrained behave in a more present-biased way. If the beneficiary withdraws her money and attempts to save at home, she has easy access to it throughout the two month period (including access to the portion she intended to save rather than spend that period); toward the end of the period she is likely to be more financially constrained and thus behave in a more present-biased way. On the contrary, if she trusts the bank and decides to save in her Bansefi account, she makes her saving decision when initially withdrawing benefits, when she is less financially constrained due to having recently received the Oportunidades payment.

It is also possible that saving money informally is difficult because the beneficiary lacks control over her husband or partner's access to money saved at home, and the husband has different (perhaps more present-biased) time preferences. Anderson and Baland (2002, p. 963) present evidence that participation in rotating savings and credit associations (ROSCAs) "is a strategy a wife employs to protect her savings against claims by her husband for immediate consumption." Consistent with differing preferences between spouses, Rubalcava et al. (2009) find that Oportunidades income (paid directly to the wife and viewed as the wife's income) tends to be spent more on investments in the future than other income does. When spouses have differing time preferences (even if neither is present-biased), the collective decision making of the household becomes present-biased (Jackson and Yariv, 2014), making soft commitment devices such as bank accounts more attractive.

We thus test whether debit cards lead to an increase in overall savings because cash saved at home is susceptible to the husband taking or requesting the money, whereas it is difficult for him to access the savings in a formal bank account.[16] Since single beneficiaries (i.e., beneficiaries who are not living with a spouse or partner) would not be affected by this barrier, a first pass to exploring

[16] On the other hand, debit cards might lead the husband to have higher access to the money, especially in households where the woman has low bargaining power. The high indirect transaction costs of a bank account without a card, and the requirement that the card holder herself appear at the bank branch to withdraw money, could make control over the husband's access to the money easier without a debit card. Indeed, in Schaner (forthcoming), women with low bargaining power are hurt by receiving debit cards because they lose control over the money. As we have already seen, however, the Bansefi bank accounts were not being used to save prior to receiving the debit cards. Even so, the beneficiary might be able to hide the money at home but unable to prevent her husband from taking and using her card to withdraw money. In our survey data, however, only 4% of beneficiaries report that their spouse sometimes withdraws money from the account.

whether beneficiaries' lack of control over their husbands' access to the money is a barrier to saving informally is to test whether a single woman responds differently to the card than a woman who is living with her spouse or partner.[17] Of course, a significant result would need to be interpreted with the caveat that single and married women have many other differences that might interact with the effect of a debit card on savings. We estimate

$$Savings_{it} = \lambda_i + \delta_t + \gamma D_{j(i)t} + \xi D_{j(i)t} \times H_i + \sum_{k \in K} \zeta_k H_i \times \mathbb{I}(t = k) + \nu_{it}, \tag{10}$$

where $Savings_{it}$ is again constructed as income minus consumption and winsorized at the 5th and 9th percentiles, H_i is a time-invariant measure of heterogeneity, and $K = \{2003, 2004, 2009\text{–}10\}$ (dropping 2002 to avoid collinearity with the household fixed effects). The $H_i \times \mathbb{I}(t = k)$ terms thus allow the evolution of savings over time to vary with H_i even in the absence of treatment. In this case, $H_i = 1$ if the beneficiary is single in the post-treatment survey wave (since marriage should not be endogenously affected by receiving the debit card and we do not want the effect to be driven by beneficiaries whose marital status changes between pre- and post-treatment).[18]

If the husband or partner's access to money is a barrier to saving for women living with a husband (or other adult) but not for single women, we expect $\gamma > 0$, $\gamma + \xi = 0$, and $\xi < 0$. Table 3 column 1 shows that we do find $\gamma > 0$ and cannot reject $\gamma + \xi = 0$, but—although the point estimate of ξ is 168 pesos (fairly close to the average treatment effect from Table A1, column 3)—it is not statistically significant from 0, so we cannot reject $\xi = 0$.

If a lack of control over the husband's access to money is indeed a barrier to saving, we would also expect treatment effect heterogeneity among women who do live with a husband or partner based on their bargaining power in the household. To test for this, we proxy for baseline female bargaining power using four questions asked only in the first wave of the survey on who makes the primary decisions in the household: whether to take their children to the doctor if they are sick, whether the children have to attend school, whether to buy them new clothes when needed, and

[17] Although it is easy to identify whether the Oportunidades beneficiary is married and living with her spouse in ENCELURB, it is difficult to determine with certainty whether unmarried beneficiaries nevertheless live with a partner. We thus include beneficiaries living in the same household as another adult who is not the household head's child or grandchild in our "non-single" group. Using this definition, the non-single group of beneficiaries is made up of 95% married women and 5% who are not married but living in the same household as another adult.

[18] The sign and statistical significance of the point estimates on ξ and γ are the same, and magnitudes similar, if we instead define H_i at baseline, where baseline refers to using the most recent pre-treatment wave in which household i was included.

"important decisions that affect the household members (transport, moving, changing jobs)." We code these questions as +1 if a woman makes the decision, 0 if spouses make them jointly, and −1 if a man makes the decision, then following Kling et al. (2007), standardize the variables to each have a mean of 0 and standard deviation of 1 and average them to create a summary measure of female bargaining power. We estimate (10) on the subset of women living with a spouse (or other adult), with H_i as this summary measure of baseline female bargaining power. Our hypothesis that women with high bargaining power could already exercise control over money saved in the home, and thus should not have as large of a treatment effect as women with low bargaining power prior to receiving the card, would mean that $\xi < 0$.

The results of this test are shown in Table 3, column 2.[19] Indeed, we find $\xi < 0$, significant at the 10% level. A one standard deviation *decrease* in baseline female bargaining power translates to an *increase* of about 198 pesos in the savings effect of the debit card, nearly as large as the average treatment effect in the full sample. This suggests that a woman with low bargaining power at baseline (and hence less control over money saved informally) receives a larger benefit from the card because it enables her to build trust in the bank and subsequently save in the account, which is out of reach of her husband. A woman with high bargaining power at baseline, on the other hand, was already able to prevent her husband from spending informal savings prior to receiving the card, and thus receives a lower benefit from the card. As a result, among women who are married or living with a partner, the savings effect caused by the debit card is higher for those with low baseline bargaining power.

Although it is common to measure bargaining power using decision-making questions similar to those used here (e.g., Ashraf et al., 2010; Antman, 2014), bargaining power can alternatively be measured based on differences in the husband and wife's age, education, literacy, and income. We now test the robustness of our result to measuring bargaining power using these variables, following the method used by Schaner (forthcoming). Specifically, we construct a Kling et al. (2007) summary measure of proxied bargaining power based on differences between spouses in these four variables at baseline. We standardize age, years of schooling, a dummy measuring literacy, and the total income earned by the beneficiary or her partner to each have a mean of 0 and standard deviation of 1. We

[19] Of the 2951 households in our sample, the Oportunidades beneficiary lives with a spouse or partner in 2098 (71%). Of these 2098, only 1625 are included in the regression for column 2; the difference of 473 households is because 93 were not included in the 2002 wave of the survey (the only wave to ask these bargaining power questions), while 380 were included but refused to answer one or more of the bargaining power questions.

then subtract the value of each variable for the husband or partner from the value for the beneficiary to create an indicator that is increasing in the woman's relative "power" in that dimension, then average across the four indicators. We again standardize the resulting bargaining power proxy so that the regression coefficient can be interpreted as the effect of a one standard deviation in female bargaining power on savings. The sample is again restricted to women who are married or living with a partner.[20] The coefficient under this alternative specification, shown in Table 3, column 3, is slightly lower than in the first specification (–163 pesos compared to –198) and is not statistically significant. Because the significance of our prior result is marginal and not robust to alternative measures of bargaining power (although, reassuringly, point estimates retain their sign and remain similar in magnitude), and because bargaining power could be correlated with unobservables that explain the heterogeneous treatment effect, we interpret the results based on differences in baseline bargaining power as merely suggestive.

Another potential barrier to saving informally is that money saved at home could be in demand from friends and relatives. It is obvious from Baland et al. (2011) and Jakiela and Ozier (2016) that the desire to conceal money in a savings account to avoid demands from others extends beyond one's spouse to friends and relatives. Ideally, we would test whether transfers from the household to other households decreased after receiving the card; the question on transfers of money to other households was not included in the post-wave survey, however. We thus estimate (10) with H_i as a dummy variable equal to 1 if the household reported giving money to other households at baseline (specifically, in any of the three pre-treatment waves). Because those with higher demands for money from friends and relatives are more likely to have $H_i = 1$, if this is a barrier to saving informally we expect $\xi > 0$. The results of this test, shown in Table 3, column 4, are inconclusive: although the point estimate on the interaction term is large, at 354 pesos, the standard error is very large, and the effect is statistically insignificant from 0. It is worth noting that only 7% of the sample has $H_i = 1$. This suggests that demands for money from relatives and friends *might* be a barrier to saving informally, but—if so—that this barrier only affects a small fraction of Oportunidades recipients.

[20]The sample includes 141 households less than in the previous specification due to households with missing values for years of schooling of the beneficiary or her partner.

9 Robustness

9.1 Internal Validity Checks

The identifying assumption for (1) and (6) is that the beneficiaries that received the debit card in waves 1 and 2 would have had the same trend in average balances and marginal propensity to save as the control group in the absence of treatment. While the assumption is inherently untestable, its plausibility was confirmed by two sets of results presented sections 2 and 4. First, although the rollout was not random, most means between treatment and control do not have statistically significant differences; there is a difference, however, in population, transfer amount, and percent of the transfer withdrawn. (For percent of the transfer withdrawn, the F-test of equality between the three means is rejected, and a test of equality of wave 1 and the control is rejected, but the test of equality between wave 2 and the control is not rejected.) More important, average balances follow parallel pre-treatment trends in wave 1 and the control prior to wave 1 receiving debit cards, and in wave 2 and the control prior to wave 2 receiving debit cards: this can be seen visually in Figure 6 and is formally tested in Section 4. The similarity of savings in the treatment and control groups before treatment contrasts sharply with the diverging trends after debit cards are received. The fact that results comparing the control to two waves receiving debit cards in different years are similar suggests this is not an artefact of a shock in a particular month or year.

Similarly, the identifying assumption for the household survey panel data results on savings, income, consumption, purchase of durables, and ownership of assets in (8) is that these variables would have followed parallel trends in the absence of treatment. Figure 14 shows these parallel trends graphically for the pre-treatment rounds of the survey. In addition, because there are many years between the last pre-treatment ENCELURB survey year (2004) and the year of treatment (2009), we supplement the ENCELURB parallel trends tests with tests using data from the 2004–2008 rounds of the ENIGH, a national income and expenditure survey used for Mexico's official poverty measurement. This is a repeated cross-section survey conducted in even years (but additionally conducted in 2005) that sampled between 20,000 and 30,000 households during each year in this time frame.

Although the publicly available version of the survey does not include each household's locality code, which determines whether the household lives in a treatment or control locality, we obtained the locality codes for sampled households from Mexico's National Institute of Statistics and Geography. Although Oportunidades receipt is reported in the survey, there is a large discrepancy between the number of beneficiaries according to the survey (after expansion factors are applied) and the number in national accounts (Scott, 2014), a problem also common in developing countries (Meyer et al., 2015), so to have sufficient power for our test we restrict the analysis to the poorest 20 percent of surveyed households to proxy for Oportunidades recipients, rather than use self-reported Oportunidades receipt. Again, the parallel trends can be clearly seen visually in Figure 14.

9.2 Alternative Explanations and Placebo Tests

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 savings effect and increasing marginal propensity to save over time. First, it could be that accumulating time *with the savings account*, rather than with the card, drives the increase over time. Second, while the hypothesis that debit cards increased trust through bank monitoring is demand-driven, the effect could be supply-driven if banks optimally responded to the increased debit card concentration by opening up more ATMs or bank branches in those localities; if such an expansion were gradual, it could explain the delayed savings effect and increasing marginal propensity to save over time. Third, the effect might be driven by locality-specific shocks unrelated to the debit cards. Fourth, the debit cards could merely make savings more salient, as in Akbas et al. (2015), by giving beneficiaries a reminder (in the form of an object carried with them) of their savings intentions.

There are a number of reasons that it is unlikely that the effects are driven by experience with the savings account leading beneficiaries to learn the benefits of saving, rather than time with the debit card itself. First, both treatment and control accounts are accumulating time with their savings accounts simultaneously. Second, because the savings accounts were mainly rolled out between 2002 and 2004 (Figure 4), most beneficiaries had already accumulated several years with the account by 2009, when treatment begins. Indeed, the median date of account opening in our 342,709 accounts is October 18, 2004, and less than 5 percent of accounts had existed for less than two years when they received debit cards. Third, our results from Section 5 include account fixed effects, so any time-invariant effect of having the account for a longer period of time would be absorbed. Fourth, to test for a time-varying effect of having the account for a longer period of time, we test whether results vary when we run the analysis separately for two groups: those who have had the account for more vs. less time. We use the median date of account opening to split the accounts into these two groups, and find that results are very similar. Appendix Figure ?? shows the equivalent of Figure 11 separately for older accounts (panels (a) and (b)), opened before the median date of October 18, 2004, and younger accounts (panels (c) and (d)) opened on or after that date.

A second possible explanation for the increase in savings over time is that banks gradually expanded complementary infrastructure in localities where treated beneficiaries live. Depending on the costs of each branch and ATM machine, this could be a profit-maximizing response to the increase in the number of debit card holders in treated localities. The increasing marginal propensity to save over time could be the result of the staggered expansion of this infrastructure, not increased trust. If this is so, then the increase in savings would have to be reinterpreted not only as the effect of debit cards but of the expansion of the whole enabling technology. Using quarterly data for each municipality on the number of bank branches and ATMs for Bansefi and all other banks, we test if there was indeed a contemporaneous expansion of infrastructure and if this was correlated geographically with Oportunidades debit card expansion or with savings in our accounts.

We first test for a relationship between the rollout of ATM cards and a supply-side expansion of banking infrastructure (ATMs and bank branches)[21] by estimating:

$$y_{mt} = \lambda_m + \delta_t + \sum_{k=-4}^{4} \beta_k D_{m,t+k} + \varepsilon_{jt},$$

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. We include one year (four quarters) of lags and one year of leads to test for a relationship between bank the debit card rollout and bank infrastructure. For this test, 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 through the last quarter of 2013 (since the rollout was from 2009 to 2012, when what we refer to as control group localities received debit cards). We separately test whether lags of credit card receipt predict banking infrastructure (i.e., whether there is a supply-side response to the rollout of debit cards) by

[21] We do not test an expansion of point of service (POS) payment terminals because the data on POS terminals by municipality does not begin until 2011, toward the end of our study period.

testing $\beta_{-4} = \cdots = \beta_{-1} = 0$ and whether leads of credit 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 $\beta_1 = \cdots = \beta_4 = 0$. We find evidence of neither relationship, failing to reject the null hypothesis of each test for each of the four dependent variables (Table 4).

To rule out locality-specific shocks that could be driving the savings effect, as opposed to the effect being driven by the debit cards, we perform a placebo test using poor non-Oportunidades households in the treated vs. control localities in the ENCELURB data. The ENCELURB initially included households deemed potentially eligible for the Oportunidades program as it was expanded to urban areas; some households did not become beneficiaries (either they were deemed ineligible or did not take up the program). Because these non-beneficiaries were "potentially eligible" for the program to be included in the survey, they are similarly (though not quite as) poor compared to the Oportunidades beneficiaries who make up our main sample. Because they did not receive debit cards during the rollout, due to not being Oportunidades beneficiaries, these individuals in treatment and control localities serve as a good placebo test for locality-level shocks. The results are presented in Figure 17a. The DID estimates on consumption, income, and savings are all insignificant from 0, although due to the low number of non-Oportunidades beneficiaries in ENCELURB (382 households), the estimates are very noisy. Nevertheless, it is comforting that the point estimates are substantially close to 0 relative to the coefficients from our main sample, and the coefficients for consumption and savings actually have the *opposite* sign as the coefficients from the main regression (shown again in panel (c) for comparison). This suggests that, although the noisy placebo estimates' 95% confidence intervals do include the point estimates from our main sample, locality-level shocks do not explain the observed results.

Finally, we test for a salience effect of the cards themselves, where the card—which a beneficiary might carry with her in a wallet or purse—serves as a salient reminder of her savings goals. In some localities, beneficiaries received their benefits through Bansefi but did not have access to a Bansefi savings account (and thus had to withdraw all of their money each pay period at a Bansefi branch); in these localities, the government decree requiring all beneficiaries to receive benefits through a plastic card led to receiving benefits on a pre-paid card, still without access to a savings account. Again using ENCELURB, we find that in localities without savings accounts that switched to a pre-paid card prior to the last round of the survey compared to localities without savings accounts that did not switch prior, there was no differential effect on consumption, income, or savings. These estimates are again noisier than the results from the main sample (here we have 2300 households), but the DID coefficient from the placebo consumption regression is statistically significant at the 10% level from the coefficient from the corresponding consumption regression in the full sample.

9.3 Alternative Barriers to Saving Informally

An alternative potential barrier to saving informally is that the money risks being stolen if saved at home. An anticipated reduction in crime was one of the primary motivations for the change to debit cards; in the U.S., changing the payment method of cash welfare payments to debit cards caused a significant decline in burglary, assault, and larceny (Wright et al., 2014). In developing countries, risk of theft has been anecdotally reported as a reason for not saving at home by cash transfer recipients in the Dominican Republic (Center for Effective Global Action, 2015), and is pointed out as a potential mechanism in Malawi by Brune et al. (2016).

To test this hypothesis, we test whether high-crime municipalities—where saving informally would be more difficult due to risk of theft—have a higher treatment effect. Specifically, we use three measures of crime at the municipal level: crimes recorded, thefts recorded, and homicides. Crimes and thefts recorded are based on reported crimes and thus suffer from reporting bias; nevertheless, total thefts is the variable that likely best represents the risks people face if saving informally. These variables come from Mexico's Executive Secretariat of the National Public Safety System. Homicides come from national vitality records and are thus less susceptible to bias. All crime rates are for the 2008 calendar year, immediately prior to the rollout of debit cards. For each of these three measures of crime, we estimate (10) where H_i (or, more precisely, $H_{m(i)}$) is either the crime rate per 100,000 in the municipality m in which household i lives (Table 5, columns 1, 3, and 5) or a dummy variable equal to 1 if the crime rate is greater than the median, with the median calculated based on the municipal crime rates faced by each beneficiary household in our sample (columns 2, 4, and 6). If risk of theft is a barrier to saving, we expect $\xi > 0$, i.e., there was a higher savings effect from the debit cards in localities with higher crime and thus greater risk that informal savings are stolen. In all specifications, the point estimate on the interaction term is statistically insignificant from 0; in many, it is close to 0; and in the majority, the sign is negative, contrary to our hypothesis. This suggests that risk of theft is not a barrier to saving informally in our context.

10 Conclusion

Although trust in financial institutions is by no means a sufficient condition to enable the poor to save, our findings suggest that it is a necessary condition. A lack of trust in banks could explain why a number of studies have found modest effects of offering savings accounts to the poor, even when these accounts have no fees or minimum balance requirements. Debit cards, a simple technology with high scale-up potential, provided beneficiaries of Mexico's large-scale cash transfer program Oportunidades with a mechanism to monitor banks by checking their balances at any bank's ATM; once beneficiaries built trust in banks, they began to save and their marginal propensity to save increased over time. We find that the observed increase in formal savings represents an increase in overall savings rather than a substitution from other forms of saving, and that beneficiaries reduce consumption of temptation goods, suggesting that saving informally is difficult and the use of financial institutions to save helps solve self-control problems.

The size of the savings effect, at 5% of income after one year with the debit card and 10% after two years, is larger than that of studies on various savings interventions such as subsidizing bank fees, increasing interest rates, and providing commitment savings devices. As a result, 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. Debit cards and other forms of mobile money, which are simple, scalable technologies that are gaining traction in many developing countries, could thus be a highly effective means of increasing financial inclusion among millions of government cash transfer recipients worldwide.

Figure 1: Low Trust in Banks Around the World

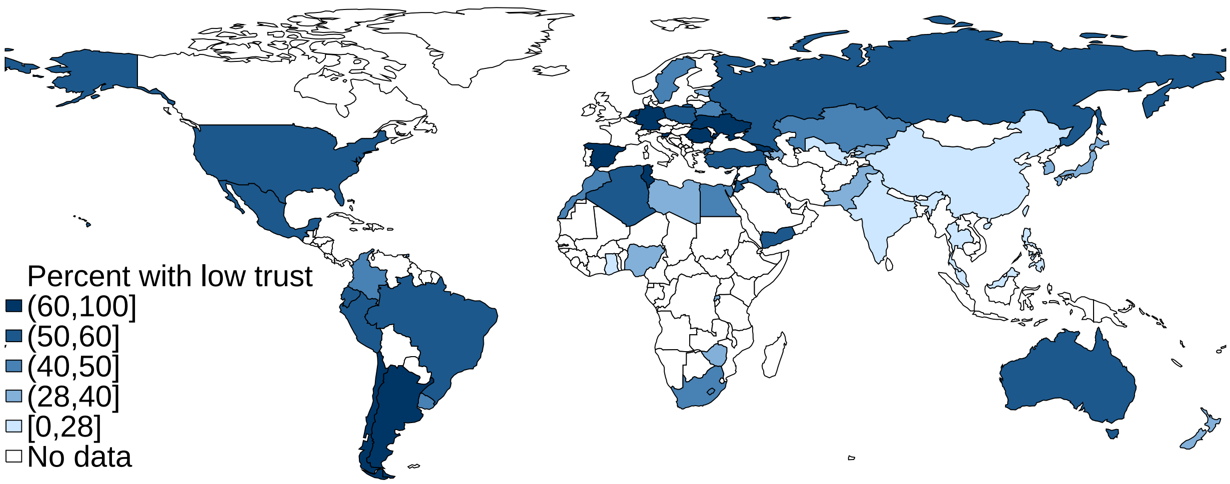

Source: World Values Survey, Wave 6 (2010–2014).
Notes: N = 82,587 individuals in 60 countries. Low trust in banks is defined as "not very much confidence" or "none at all" for the item "banks" in response to the following question: "I am going to name a number of organizations. For each one, could you 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?" Countries are divided into quintiles, with quintile cut-offs rounded to the nearest percentage point in the legend. Darker shades indicate countries with a higher percent of the population reporting low trust in banks.

Figure 2: Low Trust in Banks by Education Level in Mexico

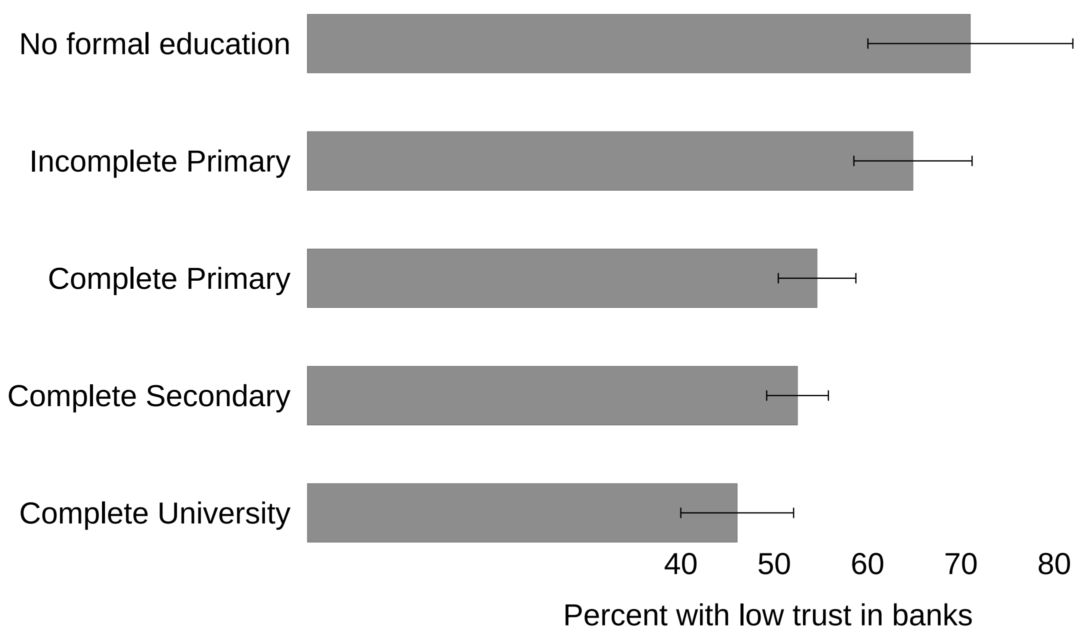

Source: World Values Survey, Mexico, Wave 6 (2012).
Notes: N = 1993 individuals. Low trust in banks is defined as "not very much confidence" or "none at all" for the item "banks" in response to the following question: "I am going to name a number of organizations. For each one, could you tell me how 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?"

Figure 3: Cross-Country Comparison of Trust in Banks and Saving in Financial Institutions

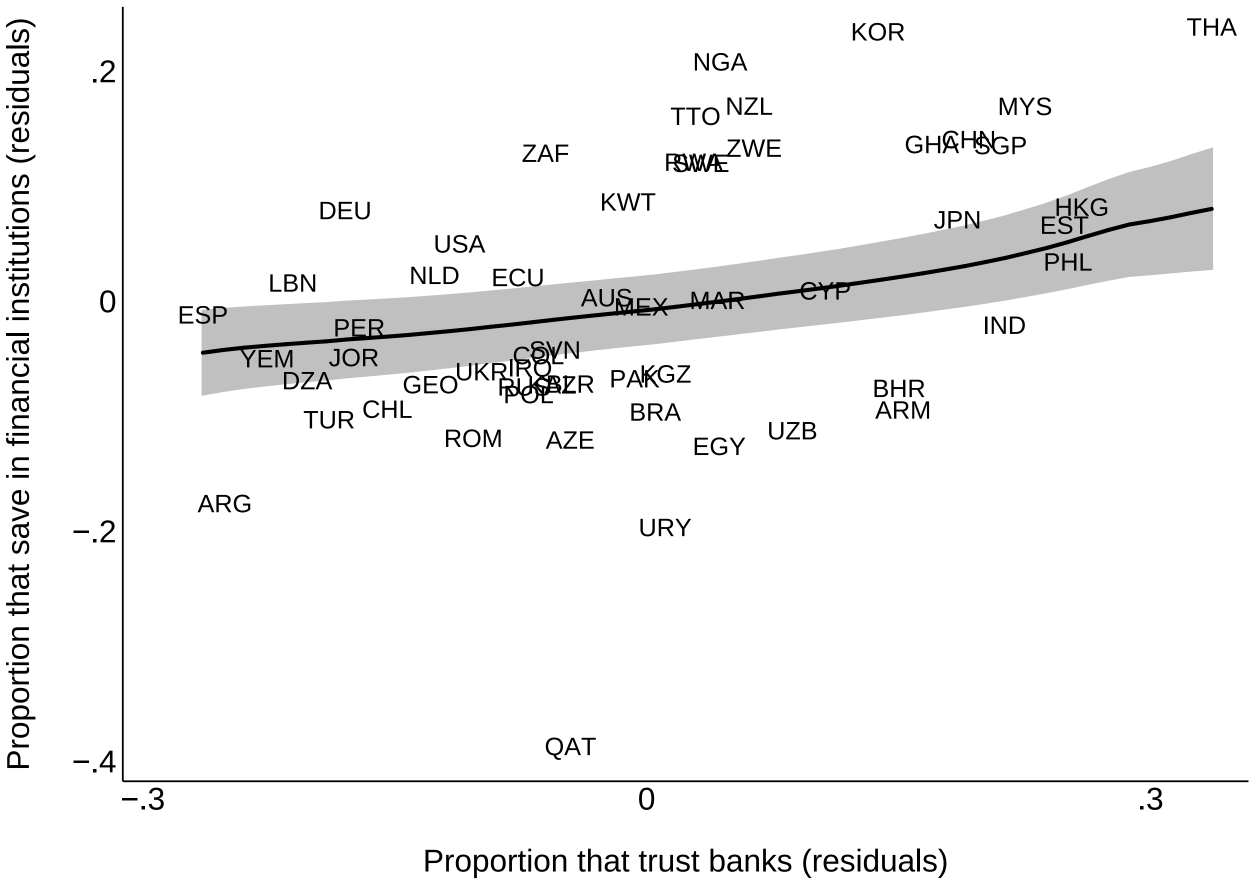

Sources: World Values Survey (WVS), Wave 6 (2010–2014); Global Findex; World Development Indicators (WDI).
Notes: $N = 56$ countries. The y-axis plots residuals from a regression of the proportion that save in financial institutions (from Global Findex) against controls (average age, education, and perceived 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 "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?" The solid line shows a kernel-weighted local polynomial regression, while the gray area shows its 95% confidence interval.

Figure 4: Timing of Roll-out and Data

(a) Administrative Bank Account Data

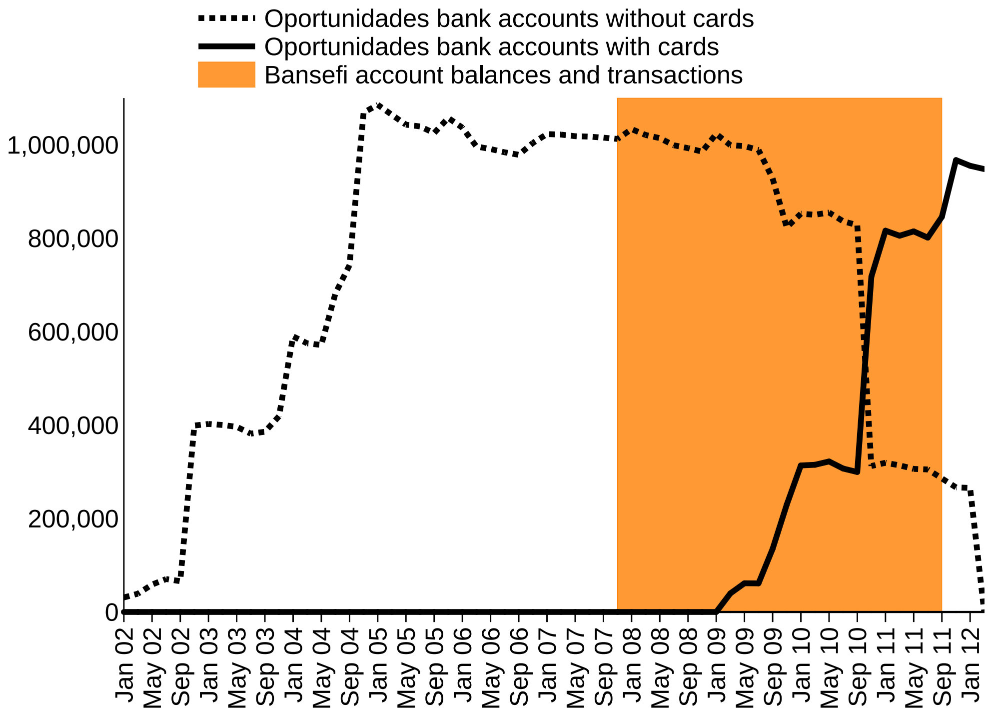

(b) Household Survey Data

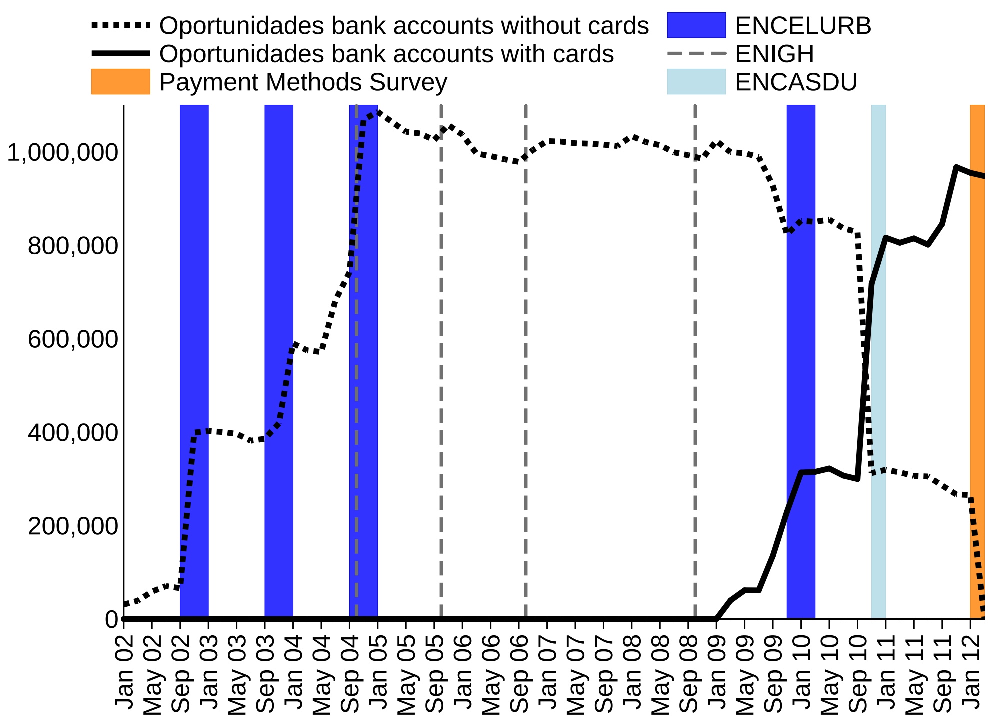

Source: Number of Oportunidades bank accounts with cards and without cards by bimester is from administrative data provided by Oportunidades.

Figure 5: Geographic Coverage and Expansion of Debit Cards

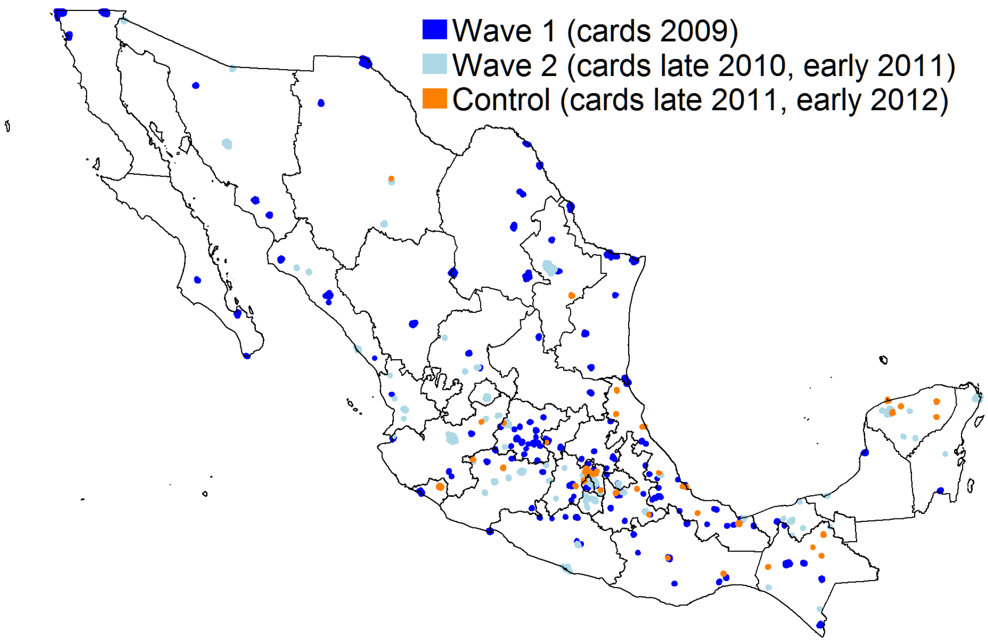

Sources: Administrative data from Oportunidades on timing of debit card receipt by locality and shape files from INEGI.
Notes: $N = 275$ localities (44 in control, 143 in wave 1, 88 in wave 2). The area of each urban locality included in the study is shaded according to its wave of treatment. Urban localities 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 6: Evolution of Average Balances

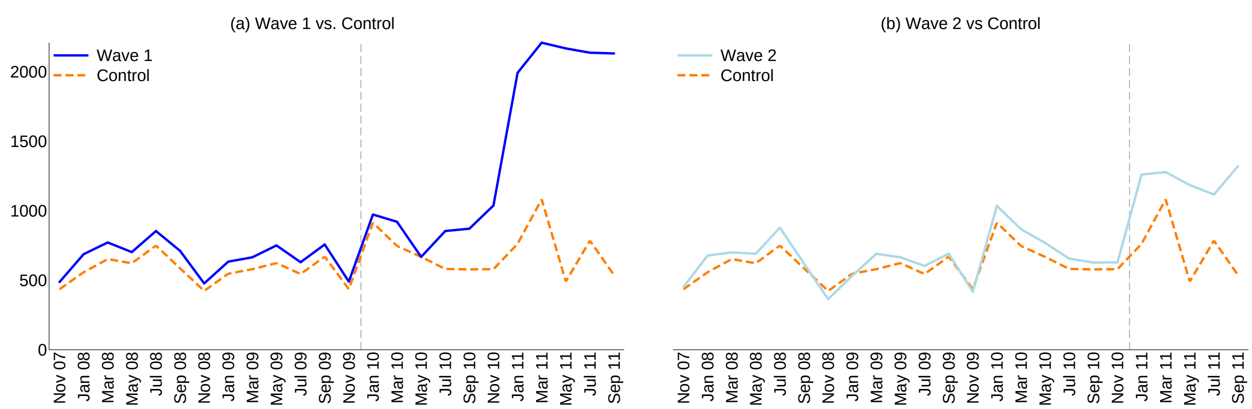

Sources: Administrative data from Bansefi on average account balances by bimester and timing of card receipt.
Notes: $N = 5{,}834{,}468$ account-bimester observations from 343,204 accounts. Average balances are winsorized at the 95th percentile.

Figure 7: Difference between Treatment and Control in Average Balances

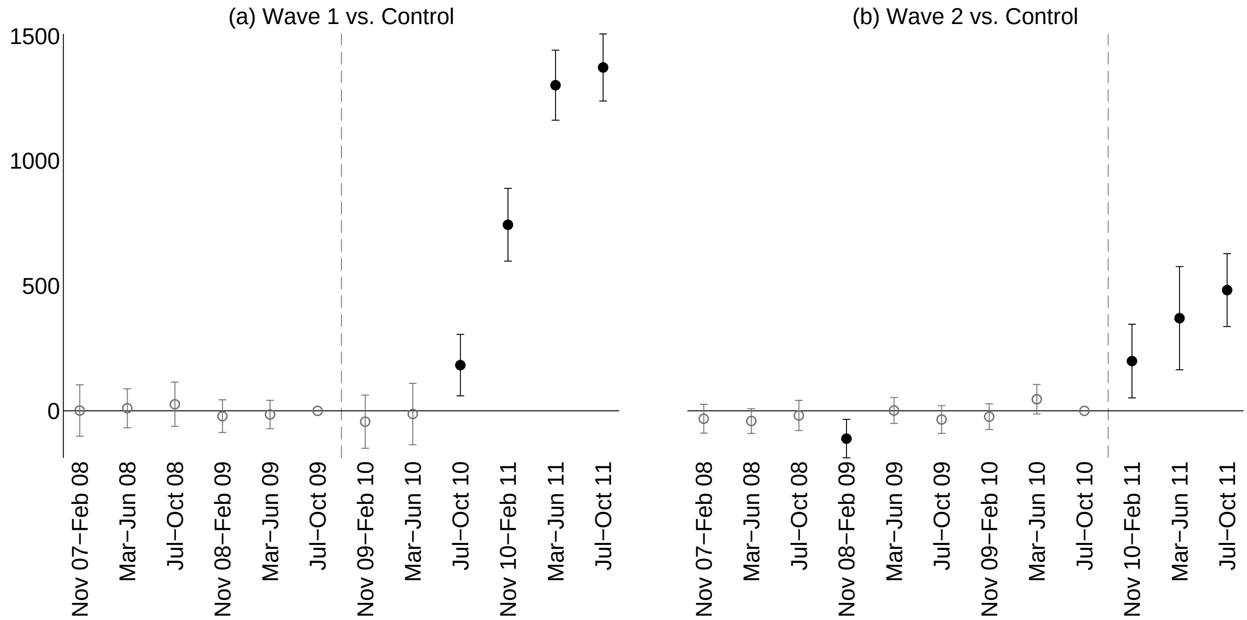

Sources: Administrative data from Bansefi on average account balances by bimester and timing of card receipt.
Notes: (a) $N = 2{,}023{,}862$ from 171,441 accounts. (b) $N = 3{,}086{,}749$ from 270,046 accounts. The figure plots ϕ_k from (1). Average balance over each four-month period is the dependent variable, and is winsorized at the 95th percentile. Whiskers denote 95 percent confidence intervals. Black filled in circles indicate results that are significant at the 5 percent level, gray filled in circles at the 10 percent 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.

Figure 8: Withdrawal to Deposit Ratio per Bimester

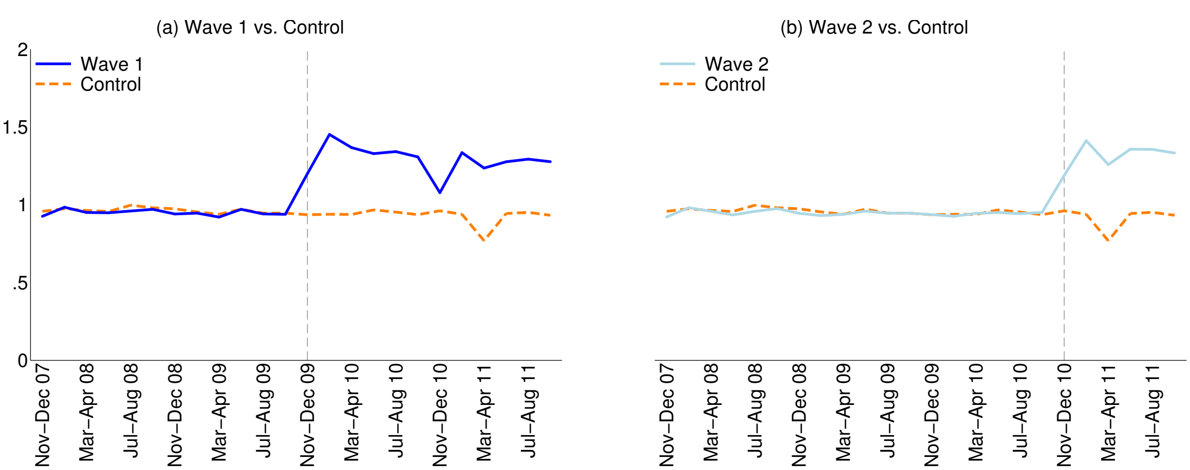

Sources: Administrative data from Bansefi on transactions by quarter and timing of card receipt.
Notes: $N = 14{,}594{,}799$ transactions from 343,204 accounts.

Figure 9: Distribution of Withdrawals

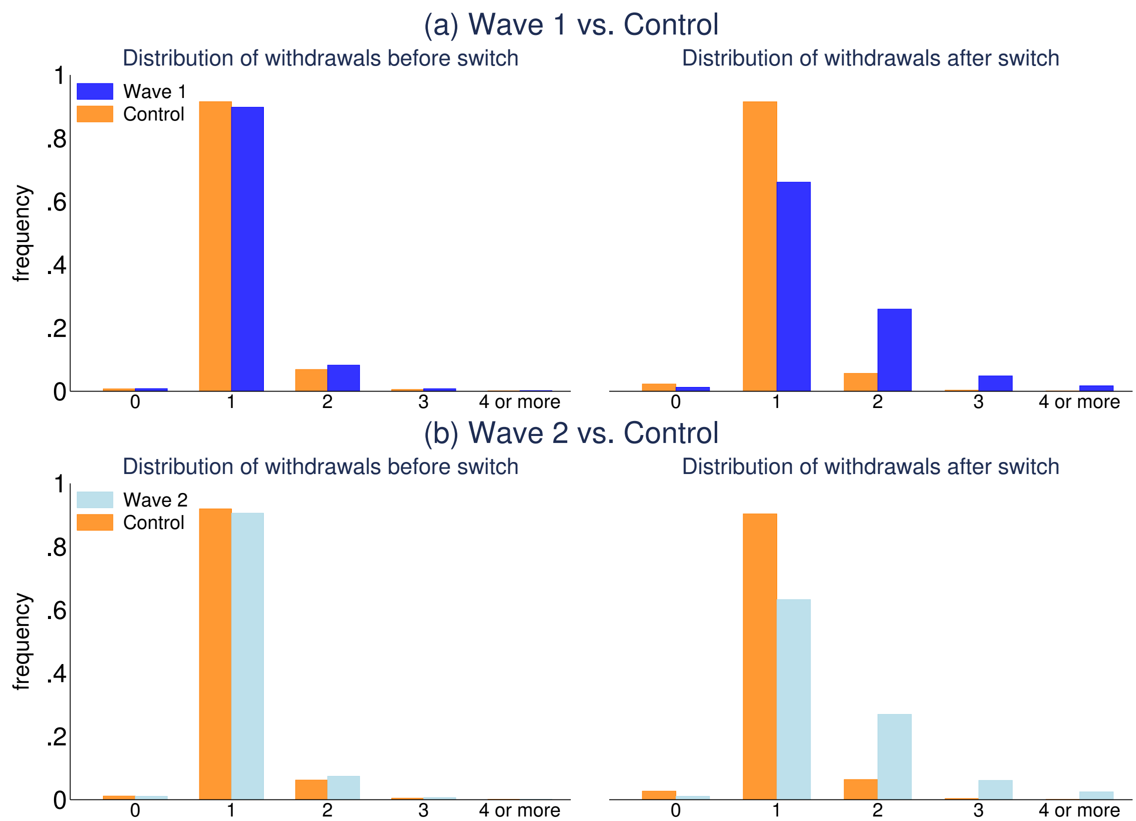

Sources: Administrative data from Bansefi on transactions by bimester and timing of card receipt.
Notes: N = 14,594,799 transactions from 343,204 accounts.

Figure 10: Difference between Treatment and Control in Net Balances

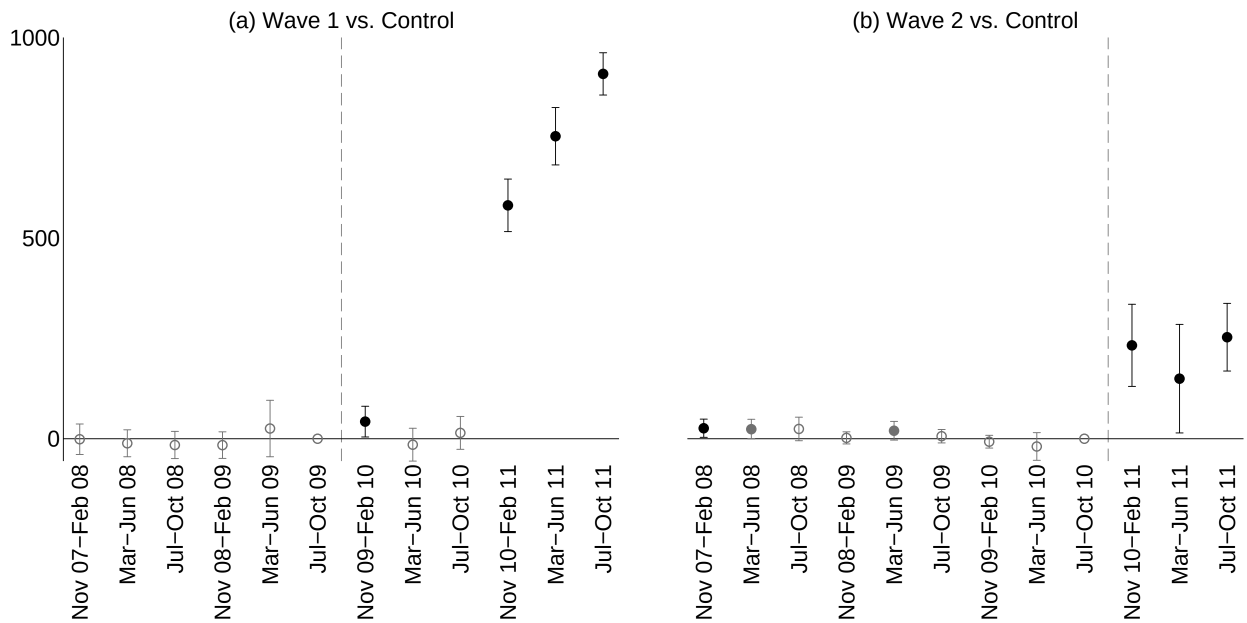

Sources: Administrative data from Bansefi on average account balances by bimester, timing and amount of transfer payments, timing and amount of withdrawals, and timing of card receipt.

Notes: (a) N = 2,023,862 from 171,441 accounts. (b) N = 3,086,749 from 270,046 accounts. 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. The figure plots ϕ_k from (1). Average balance over each four-month period is the dependent variable, and is winsorized at the 95th percentile. Whiskers denote 95 percent confidence intervals. Black filled in circles indicate results that are significant at the 5 percent level, gray filled in circles at the 10 percent 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.

Figure 11: Difference between Treatment and Control in Marginal Propensity to Save Out of Transfer

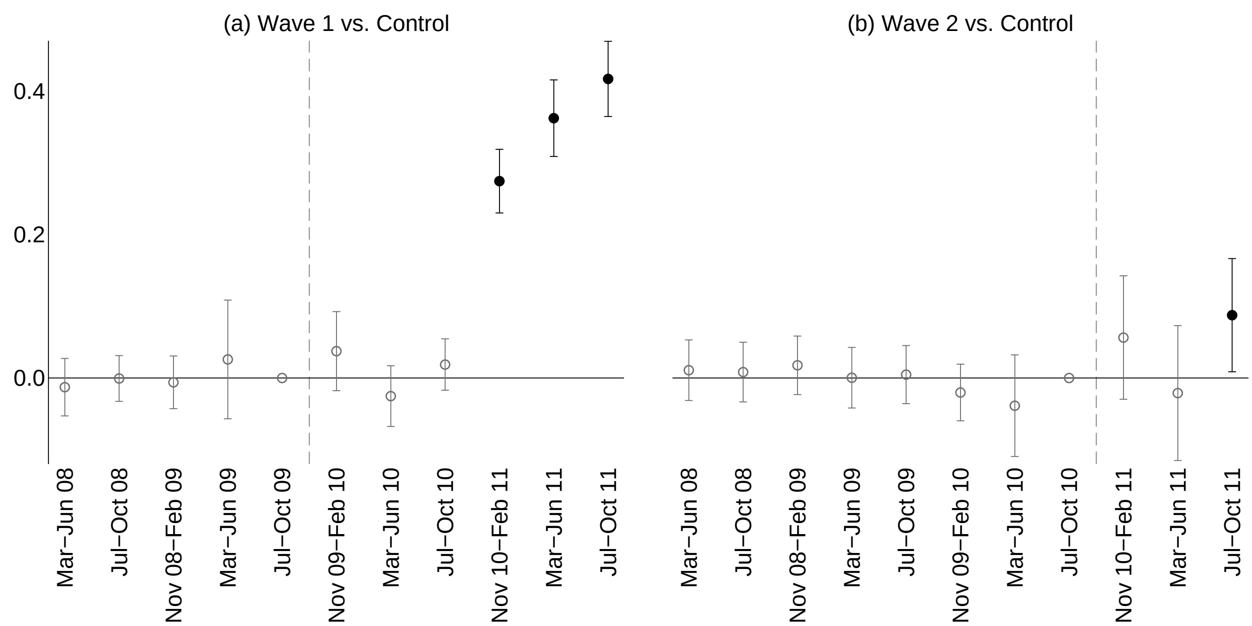

Sources: Administrative data from Bansefi on average account balances by bimester, timing and amount of transfer payments, timing and amount of withdrawals, and timing of card receipt.

Notes: (a) N = 1,852,416 from 171,441 accounts. (b) N = 2,816,671 from 270,046 accounts. (Total number of observations does not include the $t = 1$ observations, which are not included in the regressions but are used to generate $y_{ij,t-1}$ for $t = 2$ observations.) The figure plots $\alpha_k/\mu_k + \psi_k$ from (6) estimated by Blundell and Bond (1998) two-step system GMM, where μ_k is average transfers in period k. Average balances and transfer amounts are winsorized at the 95th percentile within the treatment and control groups and within each time period. The variance of $\frac{\alpha_k}{\mu_k} + \psi_k$ is estimated using the delta method. Whiskers denote 95 percent confidence intervals. Black filled in circles indicate results that are significant at the 5 percent level, gray filled in circles at the 10 percent 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.

Figure 12: Trust and Knowledge Over Time with the ATM Card

(a) ENCASDU (2010)

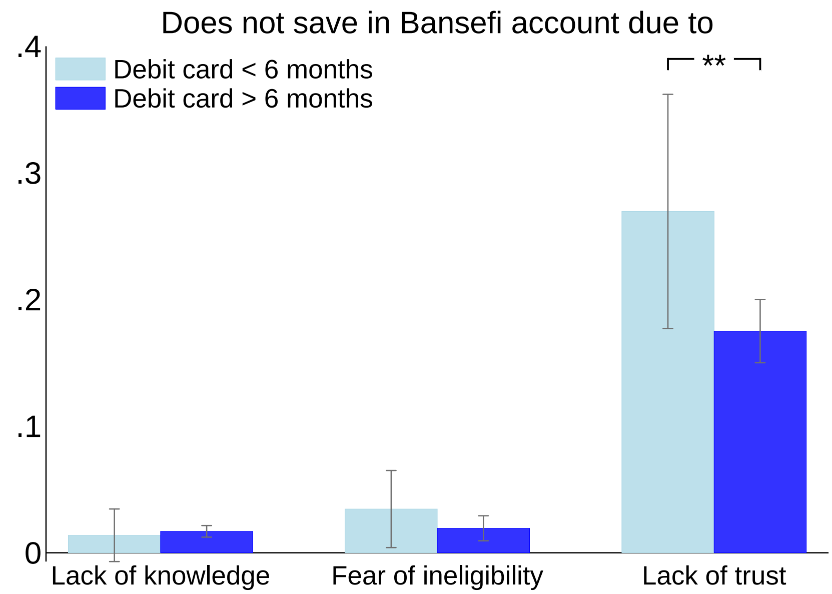

(b) Payment Methods Survey (2012)

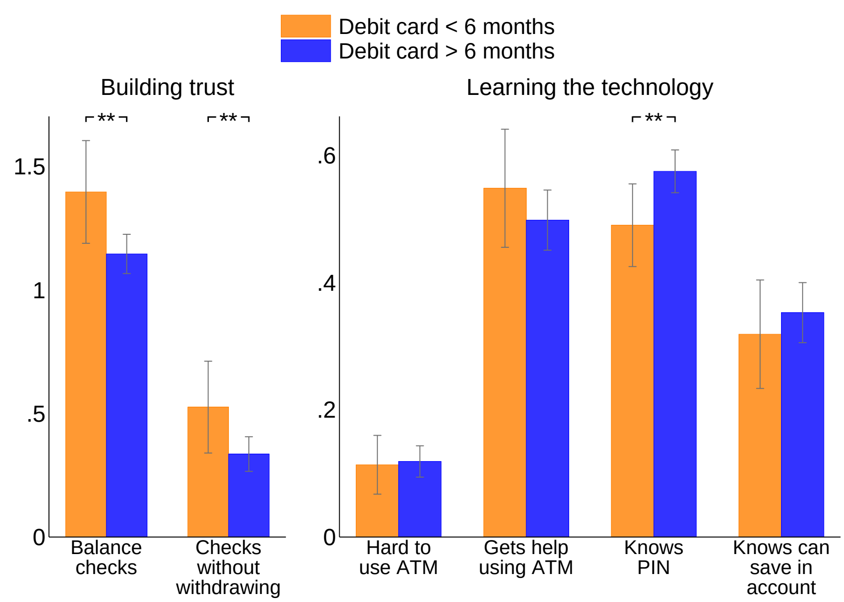

Sources: ENCASDU 2010 and Payment Methods Survey 2012.
Notes: (a) $N = 1674$. (b) $N = 1617$, or less in some regressions if there were respondents who reported "don't know" or refused to respond (see Table 2 for number of observations in each regression). Balance checks are measured over the past bimester. Whiskers denote 95 percent confidence intervals. Bars for "debit card $<$ 6 months" are colored light blue in (a) because at the time of ENCASDU 2010, those with the card 6 months or less were in wave 2 localities; bars for "debit card $<$ 6 months" are colored orange in (b) because at the time of Payment Methods Survey 2012, those with the card 6 months or less were in control localities.

Figure 13: Balance Checks (Administrative Data)

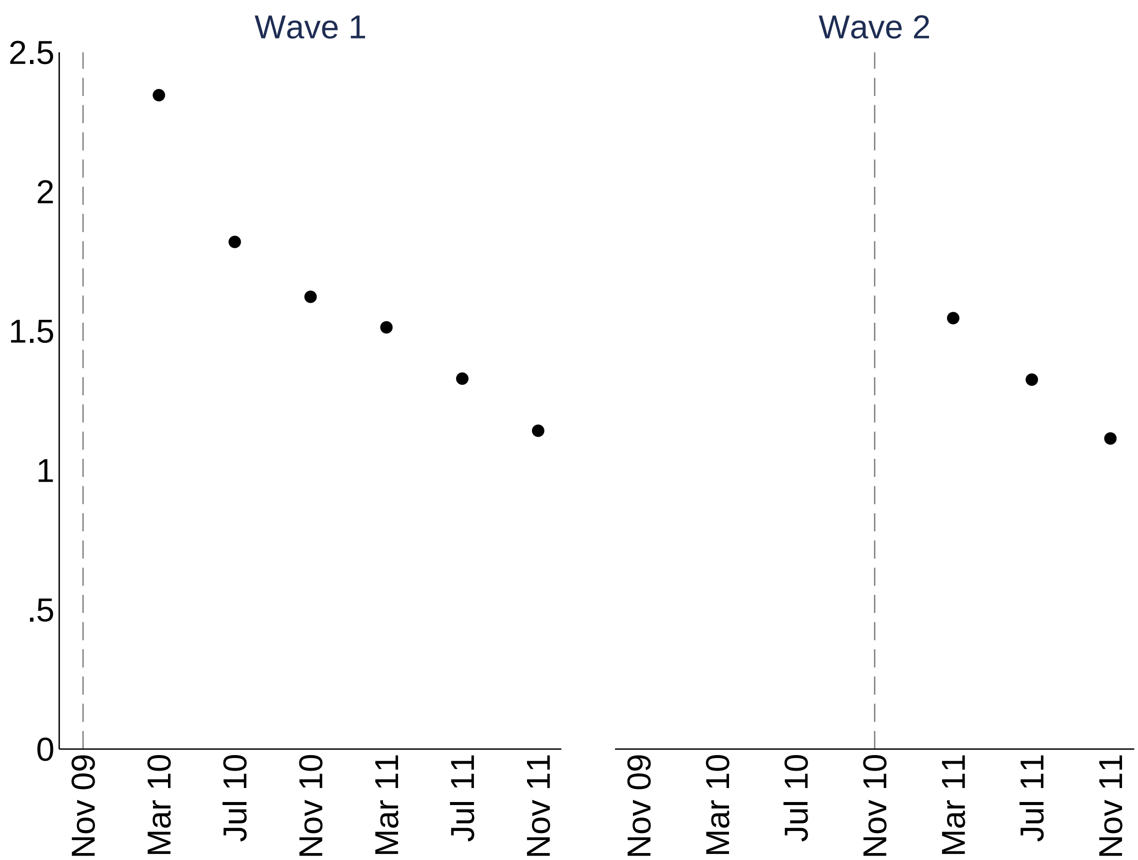

Source: Aministrative transactions data from Bansefi.
Notes: Number of balance checks per account tied to a debit card. Prior to receiving the card it was possible to check balances at Bansefi branches only, and balance checks at Bansefi branches are not recorded in our transactions data because they are free of charge.

Figure 14: Parallel Pre-Treatment Trends in Household Survey Data

(a) (b)

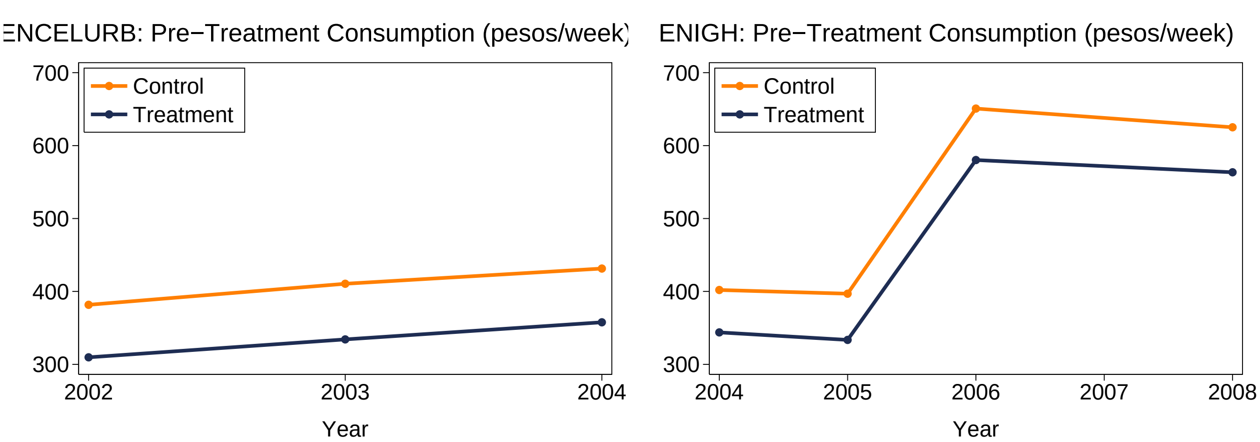

Figure 15: Effect of the debit card on consumption, income, total savings, purchase of durables, and assets

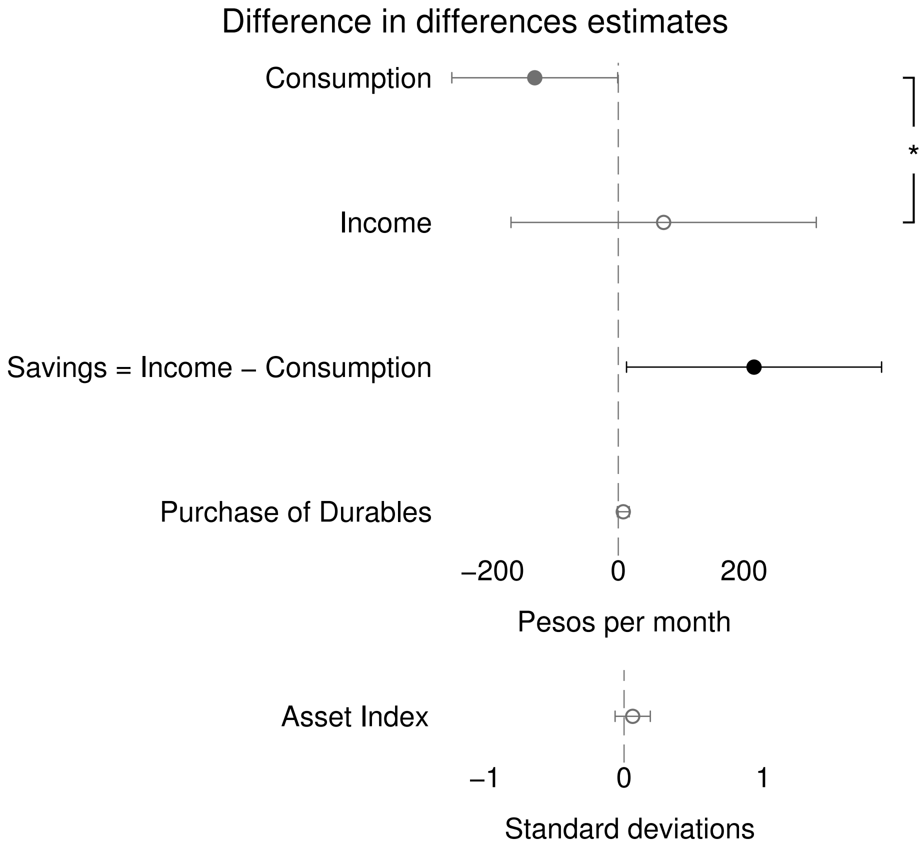

Sources: ENCELURB panel survey combined with administrative data on timing of card receipt and transfer payment histories for each surveyed beneficiary household.
Notes: $N = 11,275$ (number of households = 2951). Dependant variables are measured in pesos per month, with the exception of the asset index. 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. Whiskers denote 95 percent confidence intervals. Black filled in circles indicate results that are significant at the 5 percent level, gray filled in circles at the 10 percent 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 percent, winsorized at 5 percent, winsorized at 5 percent with baseline household characteristics interacted with time fixed effects, and winsorized at 5 percent with municipality $\times$ time fixed effects are available in Appendix Table A1. All regressions include household and time fixed effects, and standard errors are clustered at the locality level, using pre-treatment (2004) locality.

Figure 16: Effect of the debit card on different categories of consumption

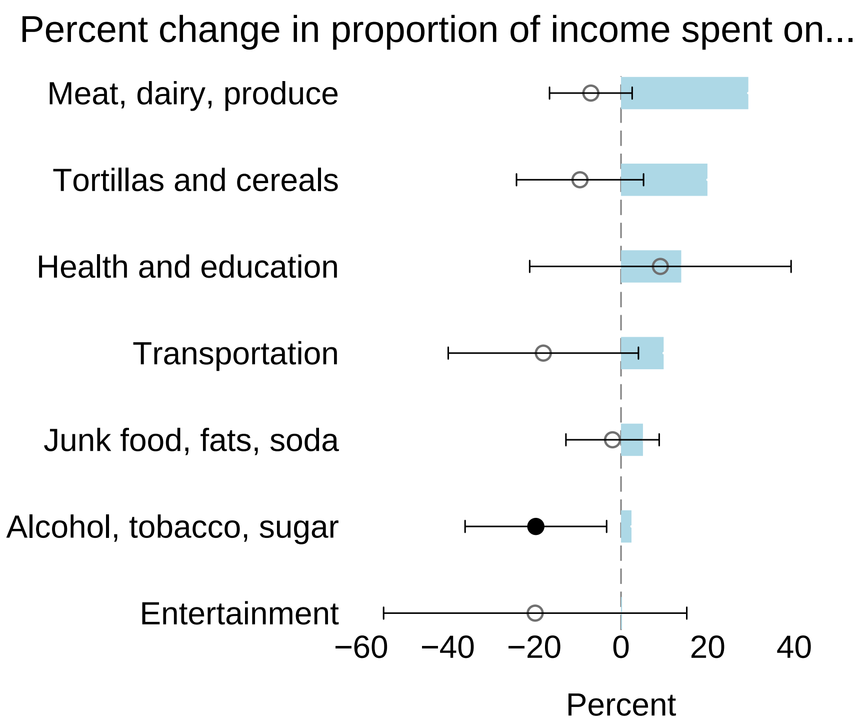

Sources: ENCELURB panel survey combined with administrative data on timing of card receipt and transfer payment histories for each surveyed beneficiary household.

Notes: $N = 11,275$ (number of households $= 2951$). Each plotted coefficient is from a separate regression using (9), and shows the percent change in the proportion of income spent on that category of consumption. In other words, the graph plots γ_g/μ_g, where μ_g is the mean proportion of income spent on consumption category g by the control group at baseline. Categories are sorted in descending order of the percent of income spent on each consumption category at baseline, i.e. $100\mu_g$, which is shown by the thick horizontal bars. The whiskers show 95% confidence intervals with no adjustment for multiple hypothesis testing. After adjusting for multiple hypothesis testing using the sharpened false discovery rate (Benjamini et al., 2006; Anderson, 2008), the result for the "alcohol, tobacco, and sugar" category is significant at the 10% rather than 5% level ($p = 0.023, q = 0.086$).

Figure 17: Placebo Tests

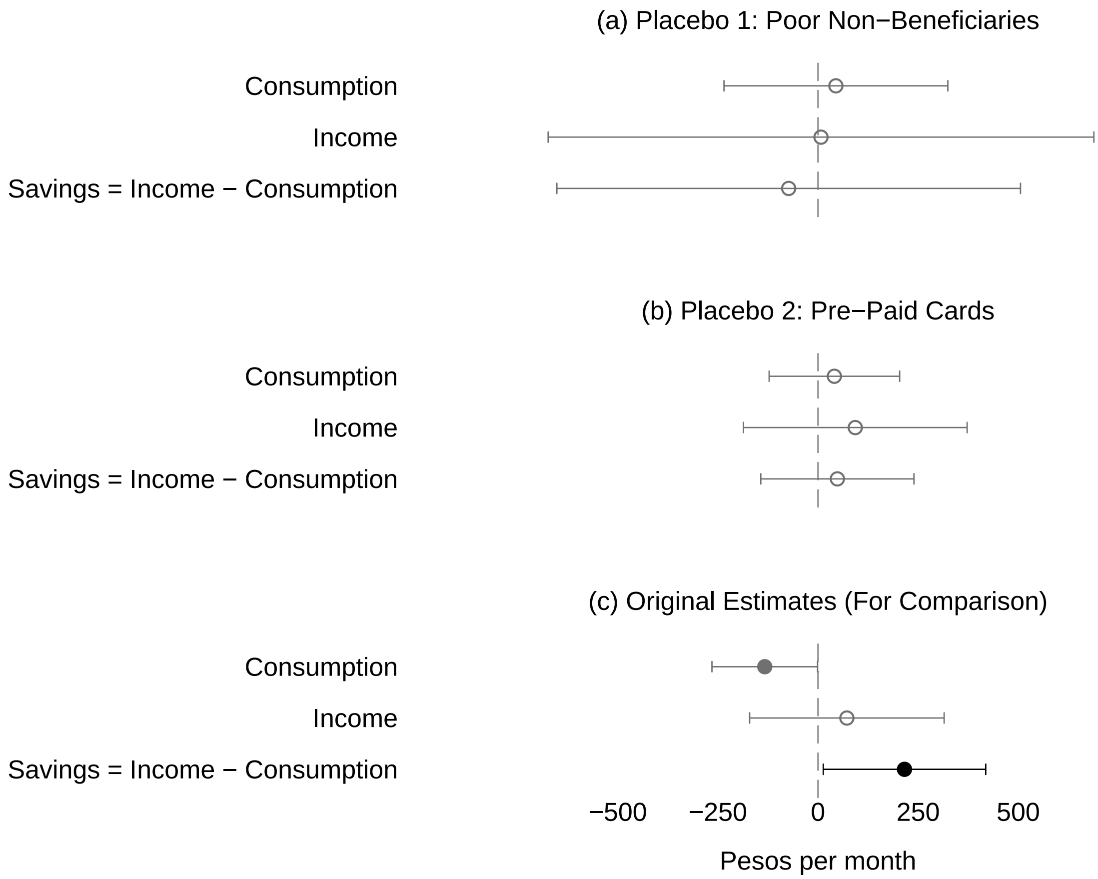

Sources: ENCELURB panel survey combined with administrative data on timing of card receipt and transfer payment histories for each surveyed beneficiary household.
Notes: (a) $N = 1415$ (number of households = 382); (b) $N = 8862$ (number of households = 2300); (c) $N = 11,275$ (number of households = 2951). Whiskers denote 95 percent confidence intervals. Black filled in circles indicate results that are significant at the 5 percent level, gray filled in circles at the 10 percent level, and hollow circles indicate results that are statistically insignificant from 0.

Table 1: Comparison of Baseline Means

Variable	Control	Wave 1	Wave 2	Diff. W1–C	Diff. W2–C	F-test p-value
Panel A: Locality-level data						
Log population	10.57	11.18	11.48	0.60***	0.91***	0.000***
	(0.11)	(0.10)	(0.16)	(0.14)	(0.19)	
Bansefi branches per 100,000	1.27	1.23	1.58	−0.03	0.32	0.411
	(0.28)	(0.13)	(0.23)	(0.30)	(0.36)	
% HHs in poverty	15.93	13.20	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 account data						
Average balance	581.25	670.32	614.29	89.07	33.05	0.112
	(12.46)	(56.24)	(21.26)	(55.33)	(23.95)	
Number of deposits	1.06	1.05	1.06	−0.02	−0.01	0.907
	(0.01)	(0.04)	(0.03)	(0.04)	(0.03)	
Size of transfer	1506.55	1809.50	1761.26	302.96***	254.71***	0.000***
	(12.73)	(20.16)	(17.47)	(23.67)	(21.15)	
Number of withdrawals	1.03	1.01	1.02	−0.01	−0.01	0.757
	(0.01)	(0.03)	(0.02)	(0.03)	(0.02)	
Percent withdrawn	98.56	97.50	99.64	−1.06**	1.08	0.021**
	(0.18)	(0.45)	(0.71)	(0.46)	(0.72)	
Years with account	5.31	5.49	5.21	0.17	−0.10	0.510
	(0.08)	(0.15)	(0.25)	(0.17)	(0.26)	
Number of accounts	97,922	73,070	171,717			

Sources: Census (2005), Bansefi branch locations (2008), poverty estimates from Oportunidades (based on 2005 Census), timing of card receipt by locality from Oportunidades, and administrative data from Bansefi.
Notes: W1 = wave 1, W2 = wave 2, C = control, Diff. = difference. For the administrative data from Bansefi, baseline is defined as January 2009 to October 2009 (prior to any accounts receiving cards in the data from Bansefi).

Table 2: Trust and Knowledge Over Time with the ATM Card

	Mean	Has card ≤ 6 months	N
Panel A: ENCASDU Survey (2010): Doesn't save in Bansefi due to ...			
Lack of knowledge	0.017***	–0.003	1,674
	(0.002)	(0.010)	
Fear of ineligibility	0.019***	0.015	1,674
	(0.004)	(0.015)	
Lack of trust	0.175***	0.095**	1,674
	(0.012)	(0.044)	
Panel B: Payment Methods Survey (2012)			
Lack of trust			
Times checked balance	1.146***	0.251**	1,493
	(0.039)	(0.105)	
Times checked balance without withdrawing	0.336***	0.190**	1,490
	(0.035)	(0.093)	
Lack of knowledge			
Hard to use ATM	0.106***	0.002	1,617
	(0.013)	(0.025)	
Gets help using ATM	0.498***	0.050	1,612
	(0.023)	(0.048)	
Knows PIN	0.575***	–0.085**	1,609
	(0.017)	(0.034)	
Knows can save in account	0.353***	–0.034	1,617
	(0.023)	(0.046)	

Notes: * indicates statistical significance at $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$. Standard errors are clustered at the locality level. The "Mean" column shows the mean for those who have had the card for more than six months; the "Has card ≤ 6 months" column shows the regression coefficient on a dummy for those who have had the debit card for six months or fewer (i.e., the difference relative to the mean column). The precise questions on trust and knowledge are as follows.

In the ENCASDU, the questions are "Do you leave part of the monetary support from Oportunidades in your bank account?" and, if the response is no, "Why don't you keep part of the monetary support from Oportunidades in your Bansefi bank account?" The regressions presented here are not conditional on saving; those who report yes to the first question are coded with trust and knowledge dependent variables of 0 and included in the regressions. The second question includes pre-written responses and an open-ended response ("other; specify"; 4% of the sample in this table responded using the open-ended option); both pre-written and open-ended responses were coded as lack of knowledge, fear of ineligibility, lack of trust, or another explanation for not saving. An example of an answer coded as lack of knowledge is "They didn't explain the process for saving." An example of an answer coded as fear of ineligibility is "Because if I save in that account they can remove me from the Oportunidades program." An example of an answer coded as lack of trust is "Because if I don't take out all the money, I can lose what remains in the bank."

In the Payment Methods Survey, each regression comes from a different survey question. These questions are: (1) Times checked balance: "In the last bimester, how many times did you consult your balance?" (2) Times checked balance without withdrawing: created by subtracting "In the last bimester, how many times did you withdraw money from the ATM?" from (1); (3) Hard to use ATM: responded "The ATM is difficult to use" (pre-written response) or a similar open-ended response to the question "What have been the main problems you have had with the ATM? [Wait for a response and record up to three of the options]"; (4) Gets help using ATM: "In general, does someone help you use the ATM?"; (5) Knows PIN: "Do you know your PIN by heart?"; (6) Knows can save in account: "Did they tell you that with the card you have a Bansefi savings account?"

Table 3: Other Barriers to Saving Informally

Dependent variable: savings	(1)	(2)	(3)	(4)
Has card at t	277.82**	241.59*	284.15*	215.15*
	(126.94)	(137.05)	(147.29)	(110.05)
Has card at t × single	–168.32			
	(176.08)			
Has card at t × baseline female bargaining power		–198.07*		
(based on self-reported decision making)		(115.56)		
Has card at t × baseline female bargaining power			–163.87	
(based on age, education, literacy, income differences)			(133.21)	
Has card at t × household gave money to others at baseline				354.82
				(419.10)
Number of households	2,951	1,625	1,484	2,951
Number of observations	11,275	6,300	5,778	11,275
Subsample	All	Not single[a]	Not single[a]	All
Time fixed effects	Yes	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes	Yes
Winsorized	5%	5%	5%	5%

Notes: [a]Not single refers to beneficiaries who live with a spouse (95% of the group) or at least one other adult (5% of the group). * 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. Dependant variable is savings, constructed as income minus consumption and measured in pesos per month. "Baseline female bargaining power" uses questions only included in the 2002 wave of the survey on who decides (i) whether to whether to take their children to the doctor if they are sick, (ii) whether the children have to attend school, (iii) whether to buy them new clothes when needed, and (iv) "important decisions that affect the household members (transport, moving, changing jobs)." The measure is constructed by coding the responses to these four questions as 1 if a woman makes the decision, 0 if they make the decision jointly, and –1 if a man makes the decision, then the responses from the multiple questions are standardized and averaged following Kling et al. (2007). "Household gave money to others at baseline" is a dummy variable equal to 1 if the household reported making transfers to others in any of the pre-treatment waves of the survey.

Table 4: Supply-Side Response

	Total		Bansefi	
	ATMs	Branches	ATMs	Branches
Current quarter	−1.52	0.03	−0.01	−0.01
	(4.14)	(0.30)	(0.01)	(0.02)
1 quarter lag	0.01	0.02	−0.02	0.02
	(4.11)	(0.34)	(0.01)	(0.02)
2 quarter lag	−10.83	0.08	0.01	0.01
	(5.64)	(0.36)	(0.03)	(0.01)
3 quarter lag	−5.42	0.08	−0.03	0.02
	(2.98)	(0.26)	(0.02)	(0.02)
4 quarter lag	−0.74	0.42	0.00	−0.03
	(5.97)	(0.50)	(0.01)	(0.03)
1 quarter lead	−1.10	−0.12	−0.01	0.00
	(3.66)	(0.36)	(0.00)	(0.02)
2 quarter lead	−6.09	0.25	0.00	0.01
	(4.90)	(0.34)	(0.02)	(0.01)
3 quarter lead	−7.84	0.25	−0.03	−0.01
	(8.00)	(0.65)	(0.01)	(0.03)
4 quarter lead	7.58	0.59	−0.01	−0.06
	(10.32)	(0.94)	(0.03)	(0.05)
Mean control group	198.29	36.87	0.49	1.41
F-test of lags	1.26	0.20	0.68	0.96
[p-value]	[0.29]	[0.94]	[0.61]	[0.43]
F-test of leads	0.69	0.44	0.79	0.67
[p-value]	[0.60]	[0.78]	[0.53]	[0.62]
Municipality fixed effects	Yes	Yes	Yes	Yes
Quarter fixed effects	Yes	Yes	Yes	Yes

Notes: * indicates statistical significance at $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$. The table shows β_k from

$$y_{jt} = \lambda_j + \delta_t + \sum_{k=-4}^{4} \beta_k D_{j,t+k} + \varepsilon_{jt}$$

where y_{jt} is the number of ATMs or bank branches of any bank or of Bansefi in municipality j during quarter t, $D_{jt} = 1$ if municipality j has at least one locality with Oportunidades debit cards in quarter t. The F-test of lags tests $\beta_{-4} = \cdots = \beta_{-1} = 0$; the F-test of leads tests $\beta_1 = \cdots = \beta_4 = 0$.

Table 5: Crime as a Barrier to Saving Informally

Dependent variable: savings	(1)	(2)	(3)	(4)	(5)	(6)
Has card at t	223.76*	189.63	151.25	153.84	295.98**	268.20*
	(117.48)	(115.45)	(117.85)	(123.30)	(116.74)	(135.26)
Has card at t × municipal crimes per 100,000	–0.07					
	(0.09)					
Has card at t × above median crimes per 100,000		–69.63				
		(206.45)				
Has card at t × municipal thefts per 100,000			–0.02			
			(0.19)			
Has card at t × above median thefts per 100,000				59.10		
				(201.25)		
Has card at t × municipal homicides per 100,000					–6.92	
					(4.46)	
Has card at t × above median homicides per 100,000						–165.77
						(217.16)
Number of households	2,951	2,951	2,951	2,951	2,951	2,951
Number of observations	11,275	11,275	11,275	11,275	11,275	11,275
Time fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Winsorized	5%	5%	5%	5%	5%	5%

Notes: * 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. Dependant variable is savings, constructed as income minus consumption and measured in pesos per month.

References

Akbas, M., Ariely, D., Robalino, D.A., Weber, M., 2015. How to help the poor to save a bit: Evidence from a field experiment in Kenya. Working Paper. URL: `http://sites.duke.edu/merveakbas/files/2014/08/How-to-help-the-poor-.pdf`.

Algan, Y., Cahuc, P., 2010. Inherited trust and growth. American Economic Review 100, 2060–2092.

Alvarez, F., Lippi, F., 2009. Financial innovation and the transactions demand for cash. Econometrica 77, 363–402.

Anderson, M.L., 2008. Multiple inference and gender differences in the effects of early intervention: A reevaluation of the abecedarian, perry preschool, and early training projects. Journal of the American Statistical Association 103, 1481–1495.

Anderson, S., Baland, J.M., 2002. The economics of roscas and intrahousehold resource allocation. Quarterly Journal of Economics 117, 963–995.

Angelucci, M., Attanasio, O., 2013. The demand for food of poor urban Mexican households: Understanding policy impacts using structural models. American Economic Journal: Economic Policy 5, 146–178.

Angelucci, M., De Giorgi, G., 2009. Indirect effects of an aid program: How do cash transfers affect ineligibles' consumption? American Economic Review 99, 486–508.

Angelucci, M., Karlan, D.S., Zinman, J., 2015. Microcredit impacts: Evidence from a randomized microcredit program placement experiment by Compartamos Banco. American Economic Journal: Applied Economics 7, 151–182.

Antman, F.M., 2014. Spousal employment and intra-household bargaining power. Applied Economics Letters 21, 560–563.

Arellano, M., Bover, O., 1995. Another look at the instrumental variable estimation of error-component models. Journal of Econometrics 68, 29–51.

Arrow, K., 1972. Gifts and exchanges. Philosophy & Public Affairs 1, 343–362.

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

Ashraf, N., Karlan, D., Yin, W., 2006. Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines. Quarterly Journal of Economics 121, 635–672.

Ashraf, N., Karlan, D., Yin, W., 2010. Female empowerment: Impact of a commitment savings product in the philippines. World Development 38, 333–344.

Ashraf, N., Karlan, D., Yin, W., 2015. Savings in transnational households: A field experiment among migrants from El Salvador. Review of Economics and Statistics 97, 332–351.

Atkinson, J., de Janvry, A., McIntosh, C., Sadoulet, E., 2013. Prompting microfinance borrowers to save: A field experiment from guatemala. Economic Development and Cultural Change 62, 21–64.

Attanasio, O., Di Maro, V., Lechene, V., Phillips, D., 2013. Welfare consequences of food prices increases: Evidence from rural Mexico. Journal of Development Economics 104, 136–151.

Attanasio, O., Pellerano, L., Reyes, S.P., 2009. Building trust? Conditional cash transfer programmes and social capital. Fiscal Studies 30, 139–177.

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

Augsburg, B., de Haas, R., Harmgart, H., Meghir, C., 2015. The impacts of microcredit: Evidence from Bosnia and Herzegovina. American Economic Journal: Applied Economics 7, 183–203.

Baland, J.M., Guirkinger, C., Mali, C., 2011. Pretending to be poor: Borrowing to escape forced solidarity in Cameroon. Economic Development and Cultural Change 60, 1–16.

Banerjee, A., Mullainathan, S., 2010. The shape of temptation: Implications for the economic lives of the poor. NBER Working Paper 15973. URL: http://www.nber.org/papers/w15973.

Banerjee, A.V., Duflo, E., Glennester, R., Kinnan, C., 2015. The miracle of microfinance? Evidence from a randomized evaluation. American Economic Journal: Applied Economics 7, 22–53.

Banfield, E.C., 1958. The Moral Basis of a Backward Society. The Free Press, Glencoe, IL.

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

Behrman, J.R., Gallardo-García, J., Parker, S.W., Todd, P.E., Vélez-Grajales, V., 2012. Are conditional cash transfers effective in urban areas? Evidence from Mexico. Education Economics 20, 233–259.

Benjamini, Y., Krieger, A.M., Yekutieli, D., 2006. Adaptive linear step-up procedures that control the false discovery rate. Biometrika 93, 491–507.

Bernheim, B.D., Ray, D., Yeltekin, S., 2015. Poverty and self-control. Econometrica 83, 1877–1911.

Blundell, R., Bond, S., 1998. Initial conditions and moment restriction sin dynamic panel data models. Journal of Econometrics 87, 115–143.

Blundell, R., Bond, S., Windmeijer, F., 2001. Estimation in dynamic panel data models: improving on the performance of the standard GMM estimators, in: Baltagi, B.H., Fomby, T.B., Hill, R.C. (Eds.), Nonstationary Panels, Panel Cointegration, and Dynamic Panels. Emerald Group Publishing, Bingley, UK. volume 15 of *Advances in Econometrics*, pp. 53–91.

Bohnet, I., Herrmann, B., Zeckhauser, R., 2010. Trust and the reference points for trustworthiness in Gulf and Western countries. Quarterly Journal of Economics 125, 811–828.

Bold, C., Porteous, D., Rotman, S., 2012. Social cash transfers and financial inclusion: Evidence from four countries. Consultative Group to Assist the Poor 77, 1–28.

Brune, L., Giné, X., Goldberg, J., Yang, D., 2016. Facilitating savings for agriculture: Field experimental evidence from malawi. Economic Development and Cultural Change 64, 187–220.

Bruno, G.S.F., 2005. Approximating the bias of the LSDV estimator for dynamic unbalanced panel data models. Economics Letters 87, 361–366.

Bryan, G., Karlan, D., Nelson, S., 2010. Commitment devices. Annual Review of Economics 2, 671–698.

Bun, M.J.G., Kiviet, J.F., 2006. The effects of dynamic feedbacks on LS and MM estimator accuracy in panel data models. Journal of Econometrics 132, 409–444.

Carvalho, L.S., Meier, S., Wang, S.W., 2016. Poverty and economic decision-making: Evidence from changes in financial resources at payday. American Economic Review 106, 260–284.

Center for Effective Global Action, 2015. Deep dive: Using pre-paid Visa cards to disburse government-to-person (G2P) payments in the Dominican Republic. Mimeo.

Chamon, M., Liu, K., Prasad, E., 2013. Income uncertainty and household savings in China. Journal of Development Economics 105, 164–177.

Chetty, R., Friedman, J.N., Leth-Petersen, S., Nielsen, T.H., Olsen, T., 2014. Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark. Quarterly Journal of Economics 129, 1141–1219.

CNBV, 2014. Comisión Nacional Bancaria y de Valores comunicado de prensa 094/2014. URL: http://www.cnbv.gob.mx/SECTORES-SUPERVISADOS/SECTOR-POPULAR/CasoFicrea/Comunicado_de_Prensa_Revocaci%C3%B3n_FICREA.pdf.

Cohn, A., Fehr, E., Maréchal, M.A., 2014. Business culture and dishonesty in the banking industry. Nature 516, 86–89.

Collins, D., Morduch, J., Rutherford, S., Ruthven, O., 2009. Portfolios of the Poor: How the World's Poor Live on $2 a Day. Princeton University Press.

Demirgüç-Kunt, A., Klapper, L., Singer, D., Oudheusden, P.V., 2015. The Global Findex Database

2014: Measuring financial inclusion around the world. Policy Research Working Ppaer 7255.
Dupas, P., Green, S., Keats, A., Robinson, J., 2016. Challenges in banking the rural poor: Evidence from Kenya's Western Province, in: Edwards, S., Johnson, S., Weil, D.N. (Eds.), Modernization and Development. University of Chicago Press, Chicago. volume 3 of *African Successes*.
Dupas, P., Robinson, J., 2013a. Savings constraints and microenterprise development: Evidence from a field experiment in Kenya. American Economic Journal: Applied Economics 5, 163–192.
Dupas, P., Robinson, J., 2013b. Why don't the poor save more? Evidence from health savings experiments. American Economic Review 103, 1138–1171.
Dynan, K.E., Skinner, J., Zeldes, S.P., 2004. Do the rich save more? Journal of Political Economy 112, 397–444.
Gambetta, D., 1988. Trust: Making and Breaking Cooperative Relations. Basil Blackwell, Oxford.
Gertler, P.J., 2004. Do conditional cash transfers improve child health? Evidence from PROGRESA's control randomized experiment. American Economic Review Papers and Proceedings 94, 336–341.
Gertler, P.J., Martinez, S.W., Rubio-Codina, M., 2012. Investing cash transfers to raise long-term living standards. American Economic Journal: Applied Economics 4, 164–192.
Guiso, L., Sapienza, P., Zingales, L., 2004. The role of social capital in financial development. American Economic Review 49, 526–556.
Guiso, L., Sapienza, P., Zingales, L., 2008. Trusting the stock market. Journal of Finance 63, 2557–2600.
Hirschman, A.O., 1984. Against parsimony: Three easy ways of complicating some categories of economic discourse. American Economic Review Papers and Proceedings 74, 89–96.
Hoddinott, J., Skoufias, E., 2004. The impact of PROGRESA on food consumption. Economic Development and Cultural Change 53, 37–61.
Iyer, R., Puri, M., 2012. Understanding bank runs: The importance of depositor-bank relationships and networks. American Economic Review 102, 1414–1445.
Jack, W., Ray, A., Suri, T., 2013. Money management by households and firms in Kenya. American Economic Review 103, 356–361.
Jack, W., Suri, T., 2014. Risk sharing and transactions costs: Evidence from Kenya's mobile money revolution. American Economic Review 104, 183–223.
Jackson, M.O., Yariv, L., 2014. Present bias and collective dynamic choice in the lab. American Economic Review 104, 4184–4204.
Jakiela, P., Ozier, O., 2016. Does Africa need a rotten kin theorem? Experimental evidence from village economies. The Review of Economic Studies 83, 231–268.
de Janvry, A., Emerick, K., Gonzalez-Navarro, M., Sadoulet, E., 2015. Delinking land rights from land use: Certification and migration in Mexico. American Economic Review 105, 3125–3149.
Karlan, D., Linden, L.L., 2014. Loose knots: Strong versus weak commitments to save for education in Uganda? mimeo.
Karlan, D., Mobius, M., Rosenblat, T., Szeidl, A., 2009. Trust and social collateral. Quarterly Journal of Economics 124, 1307–1361.
Karlan, D., Ratan, A.L., Zinman, J., 2014. Savings by and for the poor: A research review and agenda. Review of Income and Wealth 60, 36–78.
Karlan, D., Zinman, J., 2014. Price and control elasticities of demand for savings. Working Paper. URL: `http://karlan.yale.edu/sites/default/files/savingselasticities_2014_01_v9.pdf`.
Kast, F., Meier, S., Pomeranz, D., 2012. Under-savers anonymous: Evidence on self-help groups and peer pressure as a savings commitment device. Harvard Business School Working Paper 12-060. URL: `http://www.hbs.edu/faculty/Publication%20Files/12-060_`

25eb4ab1-7f56-470b-b3a9-8e039b44118e.pdf.

Kast, F., Pomeranz, D., 2014. Saving more to borrow less: Experimental evidence from access to formal savings accounts in chile. Harvard Business School Working Paper 14-001. URL: http://www.hbs.edu/faculty/Publication%20Files/14-001_2237a8d2-6147-4667-9396-cd5ae9702f03.pdf.

Kiviet, J.F., 1995. On bias, inconsistency, and efficiency of various estimators in dynamic panel data models. Journal of Econometrics 68, 53–78.

Kiviet, J.F., 1999. Expectations of expansions for estimators ina dynamic panel data model: Some results for weakly exogenous regressors, in: Hsiao, C., Lahiri, K., Lee, L.F., Pesaran, M.H. (Eds.), Analysis of panels and limited dependent variable models. Cambridge University Press, Cambridge, UK, pp. 199–225.

Kling, J.R., Liebman, J.B., Katz, L.F., 2007. Experimental analysis of neighborhood effects. Econometrica 75, 83–119.

Knack, S., Keefer, P., 1997. Does social capital have an economic payoff? a cross-country investigation. Quarterly Journal of Economics 112, 1251–1288.

La Porta, R., de Silanes, F.L., Zamarripa, G., 2003. Related lending. Quarterly Journal of Economics 118, 231–268.

Laibson, D., 2015. Why don't present-biased agents make commitments. Quarterly Journal of Economics 105, 267–272.

Levy, S., Schady, N., 2013. Latin America's social policy challenge: Education, social insurance, redistribution. Journal of Economic Perspectives 27, 193–218.

McMillan, J., Woodruff, C., 1999. Interfirm relationships and informal credit in Vietnam. Quarterly Journal of Economics 114, 1285–1320.

de Mel, S., McIntosh, C., Woodruff, C., 2013. Deposit collecting: Unbundling the role of frequency, salience, and habit formation in generating savings. American Economic Review Papers & Proceedings 103, 387–392.

Meyer, B., Mok, W.K.C., Sullivan, J.X., 2015. Household surveys in crisis. Journal of Economic Perspectives 29, 199–226.

Mullainathan, S., Shafir, E., 2009. Savings policy and decision-making in low-income households, in: Insufficient Funds: Savings, Assets, Credit and Banking Among Low-Income Households. Russell Sage Foundation Press, New York City, pp. 121–145.

Narayan, D., Pritchett, L., 1999. Cents and sociability: Household income and social capital in rural Tanzania. Economic Development and Cultural Change , 871–898.

Nickell, S., 1981. Biases in dynamic models with fixed effects. Econometrica 49, 1417–1426.

Parker, S.W., Teruel, G.M., 2005. Randomization and social program evaluation: The case of progresa. Annals of the American Acadamy of Political and Social Science 599, 199–219.

Porta, R.L., Lopez-De-Silanes, F., Shleifer, A., Vishny, R.W., 1997. Trust in large organizations. American Economic Association 87, 333–338.

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

Roodman, D., 2009a. How to do xtabond2: An introduction to difference and system GMM in Stata. Stata Journal 9, 86–136.

Roodman, D., 2009b. A note on the theme of too many instruments. Oxford Bulletin on Economics and Statistics 71, 135–158.

Rubalcava, L., Teruel, G., Thomas, D., 2009. Investments, time preferences, and public transfers paid to women. Economic Development and Cultural Change 57, 507.

Sapienza, P., Zingales, L., 2012. A trust crisis. International Review of Finance 12, 123–131.

Schaner, S., 2015. Do opposites detract? Intrahousehold preference heterogeneity and inefficient

strategic savings. American Economic Journal: Applied Economics 7, 135–174.
Schaner, S., forthcoming. The cost of convenience? Transaction costs, bargaining, and savings account use in kenya. Journal of Human Resources .
Scott, J., 2014. Redistributive impact and efficiency of Mexico's fiscal system. Public Finance Review 42, 368–390.
Suri, T., Jack, W., Stoker, T.M., 2012. Documenting the birth of a financial economy. Proceedings of the National Academy of Sciences 109, 10257–10262.
Tobin, J., 1956. The interest-elasticity of transactions demand for cash. Review of Economics and Statistics 38, 241–247.
Williamson, O.E., 1993. Calculativeness, trust, and economic organization. Journal of Law and Economics 36, 453–486.
Windmeijer, F., 2005. A finite sample correction for the variance of linear efficient two-step GMM estimators. Journal of Econometrics 126, 25–51.
Wright, R., Tekin, E., Topalli, V., McClellan, C., Dickinson, T., Rosenfeld, R., 2014. Less cash, less crime: Evidence from the Electronic Benefit Transfer program. NBER Working Paper 19996. URL: http://www.nber.org/papers/w19996.
Ziliak, J.P., 1997. Efficient estimation with panel data when instruments are predetermined: An empirical comparison of moment-condition estimators. Journal of Business and Economic Statistics 15, 419–431.
Zinman, J., 2009. Credit or debit? Journal of Banking and Finance 33.

Appendix A: Additional Figures and Tables

Figure A1: Difference between Treatment and Control in Marginal Propensity to Save Out of Transfer, Separated by Time with Account

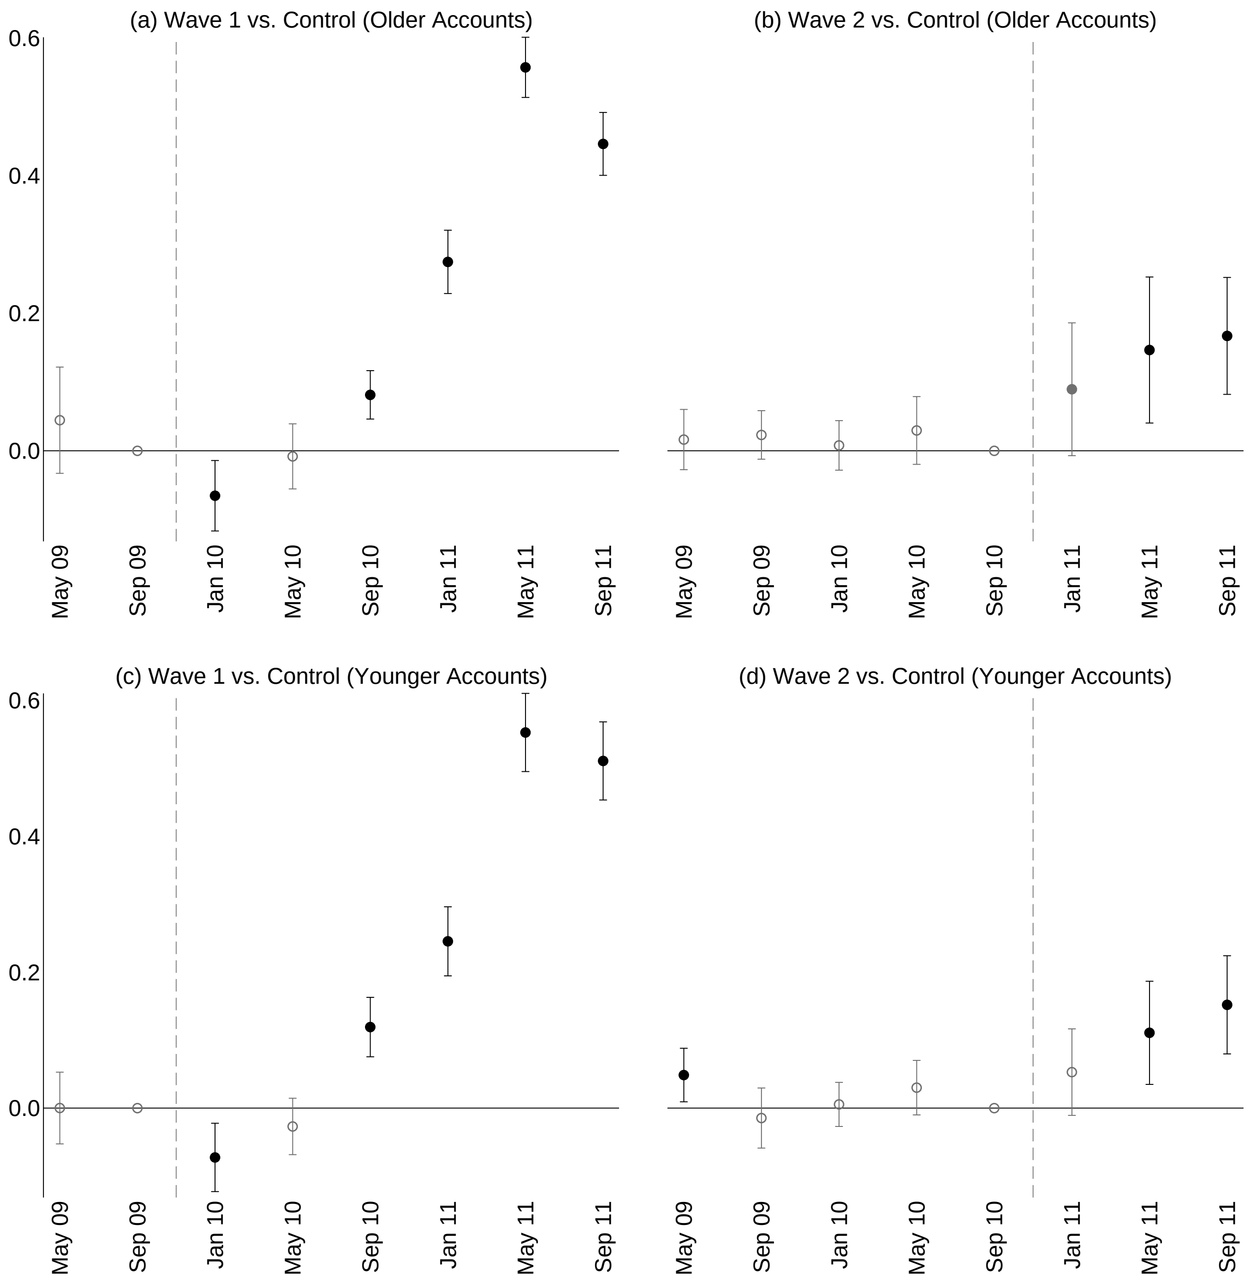

Sources: Administrative data from Bansefi on average account balances by bimester, transfer payments, and timing of card receipt.
Notes: (a) N = 743,776 from 99,362 accounts; (b) N = 905,335 from 118,228 accounts; (c) N = 455,172 from 79,511 accounts; (d) N = 1,088,677 from 157,717 accounts. See the notes to Figure 11 for the specification. Accounts are split into older accounts and younger accounts based on the median account opening date, which is October 18, 2004.

Table A1: Change in Savings and Assets After Receiving Card

	(1)	(2)	(3)	(4)	(5)
Consumption	–178.76**	–144.01*	–132.90*	–251.28**	–149.74**
	(85.60)	(74.76)	(67.30)	(116.00)	(68.58)
Income	85.91	106.55	72.16	40.93	69.71
	(158.58)	(138.41)	(123.88)	(137.48)	(122.05)
P-value Consumption vs. Income	[0.055]*	[0.033]**	[0.071]*	[0.008]***	[0.053]*
Savings = Income – Consumption	264.66*	241.30**	215.98**	285.42**	234.59**
	(134.64)	(114.94)	(103.50)	(118.95)	(104.74)
Purchase of durables	5.94	6.22	7.99	6.55	6.91
	(12.55)	(8.52)	(4.82)	(6.78)	(4.55)
Asset index	0.04	0.04	0.06	–0.07	0.05
	(0.07)	(0.07)	(0.06)	(0.06)	(0.06)
Number of households	2,951	2,951	2,951	2,951	2,938
Number of observations	11,275	11,275	11,275	11,275	11,243
Time fixed effects	Yes	Yes	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes	Yes	Yes
Municipality × time fixed effects	No	No	No	Yes	No
Household characteristics × time	No	No	No	No	Yes
Winsorized	No	1%	5%	5%	5%

Notes: * 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. Dependant 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 (5) is slightly lower because 13 households have missing values for one of the household characteristics included (interacted with time fixed effects) in that specification.

Appendix B: Mechanical Effect

This appendix explains the computation of the mechanical effect for every pattern of deposit and withdrawal that occurs in a bimester. The mechanical effect is defined as the contribution to average balances from the transit of the Opportunidades transfers on recipients' accounts, and interpreted as the part of average balances that does not represent net savings in the bimester.

Table B1: Computation of Mechanical Effect

	Pattern	% Total	Conditions	Mechanical Effect
Panel A. Regular patterns: single deposit into account in the bimester				
(1)	DW	73.4%	$W \leq D$	$lapse_{DW} * Withdrawal$
			$W > D$	$lapse_{DW} * Deposit$
(2)	DWW	9.1%	$W_1 \geq D$	$lapse_{DW_1} * Deposit$
			$W_1 < D$ & $W_1 + W_2 \geq D$	$lapse_{DW_1} * Withdrawal_1 + lapse_{DW_2} * (Deposit - Withdrawal_1)$
			$W_1 + W_2 < D$	$lapse_{DW_1} * Withdrawal_1 + lapse_{DW_2} * (Withdrawal_2)$
(3)	DWWW	1.7%	$W_1 \geq D$	$lapse_{DW_1} * Deposit$
			$W_1 < D$ & $W_1 + W_2 \geq D$	$lapse_{DW_1} * Withdrawal_1 + lapse_{DW_2} * (Deposit - Withdrawal_1)$
			$W_1 + W_2 < D$ & $W_1 + W_2 + W_3 \geq D$	$lapse_{DW_1} * Withdrawal_1 + lapse_{DW_2} * Withdrawal_2$
				$+lapse_{DW_3} * (Deposit - Withdrawal_1 - Withdrawal_2)$
Panel B. Irregular patterns: multiple deposits into account in the bimester				
(4)	DDWW	3.1%	$W_1 \leq D_1$ & $W_2 \leq D_2$	$lapse_{D_1W_1} * Withdrawal_1 + lapse_{D_2W_2} * Withdrawal_2$
			$W_1 > D_1$ & $W_2 \leq D_2$	$lapse_{D_1W_1} * Deposit_1 + lapse_{D_2W_2} * Withdrawal_2$
			$W_1 \leq D_1$ & $W_2 < D_2$	$lapse_{D_1W_1} * Withdrawal_1 + lapse_{D_2W_2} * Deposit_2$
			$W_1 > D_1$ & $W_2 > D_2$	$lapse_{D_1W_1} * Deposit_1 + lapse_{D_2W_2} * Deposit_2$
(5)	DWD	3.0%	$W \leq D$	$lapse_{DW} * Withdrawal$
			$W > D$	$lapse_{DW} * Deposit$
(6)	DDW	2.7%	$W \geq D_1 + D_2$	$lapse_{D_1W} * Deposit_1 + lapse_{D_2W} * Deposit_2$
			$W < D_1 + D_2$ & $W \leq D_2$	$lapse_{D_1W} * (Withdrawal - Deposit_2) + lapse_{D_2W} * Deposit_2$
			$W < D_2$	$lapse_{D_2W} * Withdrawal$
(7)	DWDW	1.6%	$W_1 \leq D_1$ & $W_2 \leq D_2$	$lapse_{D_1W_1} * Withdrawal_1 + lapse_{D_2W_2} * Withdrawal_2$
			$W_1 > D_1$ & $W_2 \leq D_2$	$lapse_{D_1W_1} * Deposit_1 + lapse_{D_2W_2} * Withdrawal_2$
			$W_1 \leq D_1$ & $W_2 < D_2$	$lapse_{D_1W_1} * Withdrawal_1 + lapse_{D_2W_2} * Deposit_2$
			$W_1 > D_1$ & $W_2 > D_2$	$lapse_{D_1W_1} * Deposit_1 + lapse_{D_2W_2} * Deposit_2$

D_i indicate the i^{th} deposit and W_i indicate 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 (up to the first four transactions per bimester of) all patterns have been coded to obtain an estimate of the mechanical effect.