

IDB WORKING PAPER SERIES N° (IDB-WP-01018)

Increasing Retirement Savings through Access Points and Persuasive Messages: Evidence from Mexico

Mariano Bosch
Adrian Rubli

Inter-American Development Bank
Labor Markets Division

April 2019

Increasing Retirement Savings through Access Points and Persuasive Messages: Evidence from Mexico

Mariano Bosch
Adrian Rubli

Cataloging-in-Publication data provided by the
Inter-American Development Bank
Felipe Herrera Library
Bosch, Mariano.

Increasing retirement savings through access points and persuasive messages:
evidence from Mexico / Mariano Bosch, Adrian Rubli.

p. cm. — (IDB Working Paper Series ; 1018)

Includes bibliographic references.

1. Retirement income-Mexico-Econometric models. 2. Pensions-Mexico-Econometric models. 3. Transaction costs-Mexico-Econometric models. I. Rubli, Adrian. II. Inter-American Development Bank. Labor Markets Division. III. Title. IV. Series.

IDB-WP-1018

<http://www.iadb.org>

Copyright © 2019 Inter-American Development Bank. This work is licensed under a Creative Commons IGO 3.0 Attribution-NonCommercial-NoDerivatives (CC-IGO BY-NC-ND 3.0 IGO) license (<http://creativecommons.org/licenses/by-nc-nd/3.0/igo/legalcode>) and may be reproduced with attribution to the IDB and for any non-commercial purpose, as provided below. No derivative work is allowed.

Any dispute related to the use of the works of the IDB that cannot be settled amicably shall be submitted to arbitration pursuant to the UNCITRAL rules. The use of the IDB's name for any purpose other than for attribution, and the use of IDB's logo shall be subject to a separate written license agreement between the IDB and the user and is not authorized as part of this CC-IGO license.

Following a peer review process, and with previous written consent by the Inter-American Development Bank (IDB), a revised version of this work may also be reproduced in any academic journal, including those indexed by the American Economic Association's EconLit, provided that the IDB is credited and that the author(s) receive no income from the publication. Therefore, the restriction to receive income from such publication shall only extend to the publication's author(s). With regard to such restriction, in case of any inconsistency between the Creative Commons IGO 3.0 Attribution-NonCommercial-NoDerivatives license and these statements, the latter shall prevail.

Note that link provided above includes additional terms and conditions of the license.

The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the Inter-American Development Bank, its Board of Directors, or the countries they represent.



Increasing Retirement Savings through Access Points and Persuasive Messages: Evidence from Mexico*

Mariano Bosch[†]
IDB

Adrian Rubli[‡]
ITAM

April 2019

Transaction costs impose a barrier to savings, but lowering them may have smaller impacts than expected due to other constraints, such as psychological biases. Within the context of retirement savings under defined contributions in the privatized pension system in Mexico, we analyze two staggered policies: an expansion of access channels for additional contributions through 7-Eleven stores, and the bundling of this policy with a media campaign providing persuasive reminders to save. We find that access alone has small and weak effects on savings, but including the persuasive message leads to a 10-12% increase in savings behavior that is long-lasting.

JEL codes: D14, D83, G23

Key words: retirement savings, access, transaction costs, media campaigns, non-branded advertising, persuasion

*For their helpful comments, we are grateful to Andrew Foster, Paolina Medina, and seminar participants at ITAM and the 2018 Latin American and Caribbean Economic Association (LACEA). We thank the staff at CONSAR, especially Alejandro Maldonado, Alejandra Vargas, Marlene Olvera, Iveth García, and Fernanda Santiago, for continued support in accessing the data and providing additional information. We also thank Anne Hand and María Teresa Silva at IDB for their constant support and helpful comments. Adrian Rubli acknowledges support from the Asociación Mexicana de Cultura, AC. The views expressed herein are those of the authors and do not necessarily reflect the official views of CONSAR or any other government or non-government agency or private enterprise. All remaining errors are our own.

[†]Inter-American Development Bank, Labor Markets and Social Security Unit. 1300 New York Ave., NW Washington, DC, USA. Email: *mbosch@iadb.org*.

[‡]Corresponding author. Instituto Tecnológico Autónomo de México, Department of Business Administration. Camino a Santa Teresa 930. Colonia Heroes de Padierna, Magdalena Contreras CDMX 10700, Mexico. Email: *adrian.rubli@itam.mx*.

1 Introduction

There is a large literature concerned with low savings rates across different contexts (Karlan and Morduch, 2010).¹ In particular for retirement savings, for the 64% of American families with 401(k) accounts, median savings is just 60,000 USD (Morrissey, 2016). In Mexico – the setting for this study – income replacement rates at retirement will be on average around 40% (Villagómez, 2014).² For individuals with access to financial services, transaction costs, both monetary and non-pecuniary, and psychological barriers, such as self-control and time inconsistency, may impose important constraints on savings behavior.³

Studies have documented the causal impact of reducing transaction costs, for example through debit cards and access to ATMs (Schaner, 2017; Bachas et al., 2018). The literature has also shown the role that psychological biases play, for example, by analyzing commitment devices (Ashraf et al., 2006; Dupas and Robinson, 2013). However, little is known about the interaction of the two, especially in terms of the relative efficiency of bundling policies that reduce both transaction costs and psychological barriers.

This paper aims to fill that gap by analyzing two policies within the context of voluntary contributions to the privatized social security system in Mexico. We analyze a first policy that decreased transaction costs by making deposits available at 7-Eleven convenience stores (we call this an increase in access channels), and a subsequent ad campaign that provided a non-informative, persuasive reminder to save (we call this the bundled policies of both access expansion and media campaign).

Life-cycle fluctuations in earnings are an important motivation for savings (Karlan and Morduch, 2010). Defined contribution plans for retirement are becoming increasingly popular, especially as more economies transition to privatized pension systems (Whitehouse,

¹It is not ex-ante clear whether *everyone* should be saving more (Karlan et al., 2014). However, policy-makers and academics have pushed towards reducing savings constraints in an effort to increase savings.

²Financial advisors generally recommend a replacement rate of roughly 70%. See, for example, Biggs and Springstead (2008).

³The literature has identified multiple other constraints to saving, such as access, loss aversion, mental accounting, financial literacy, advertising, peer effects, and technology (see Karlan and Morduch, 2010, for a review).

2006; Pallares-Miralles et al., 2012). These individual retirement accounts allow for more flexibility in investment decisions (Benartzi and Thaler, 2013; Krasnokutskaya and Todd, 2009). If workers are rational agents maximizing lifetime utility, this flexibility should allow adjustments to shocks. However, frictions and mental gaps often lead to individuals using rules of thumb instead of optimal behavior (Handel and Schwartzstein, 2018).⁴

In 1997, Mexico introduced a new privatized pension system based on defined contribution plans. Formal-sector workers are automatically enrolled and allowed to choose a fund manager. There is a mandatory 6.5% contribution rate, with employers, workers, and the federal government all contributing a share. Workers may also choose to increase their contributions through additional deposits called voluntary savings.

Voluntary contributions are kept separate from the main contributions, and workers are allowed to make withdrawals before retirement. Considering Mexico’s low financial inclusion, both in terms of access to banking services as well as interest rates that are often below inflation, these voluntary contributions can effectively be used for non-retirement savings, and actually provide the highest financial return for the majority of the population.

Prior to the policies we analyze, workers could make voluntary contributions either by requesting direct payroll deductions from their employer or by directly contacting their fund manager. To facilitate access, the government regulatory agency overseeing the pension system, CONSAR (National Commission of the Retirement Savings System), partnered with 7-Eleven to allow workers to make voluntary deposits directly at these convenience stores with as little as 50 pesos (3.70 USD) by simply providing their national id number.

Several months later, CONSAR implemented a nation-wide, non-branded TV and radio advertising campaign, which ran continuously for six months. The ads emphasized that saving for retirement, even in small quantities, is important, and depicted workers mak-

⁴The literature has analyzed various heuristics, such as inertia (Chetty et al., 2014; Choi et al., 2003, 2005; Illanes, 2016; Luco, 2013); loss aversion (Benartzi and Thaler, 1995; Looney and Hardin, 2009); peer effects (Duflo and Saez, 2002); framing of information and complexity of enrollment and investment procedures (Carroll et al., 2009; Iyengar et al., 2004; Hastings and Tejada-Ashton, 2008); and even biases driven by grammatical features of language (Chen, 2013).

ing contributions at 7-Eleven locations. The ads, however, did not provide any financial information.

The first intervention alleviates transaction costs by increasing access channels for deposits. The bundle adds a persuasive message that addresses psychological biases like forgetfulness and procrastination, and that may also provide information on the access expansion. We exploit 7-Eleven’s plausibly exogenous market presence, the staggered rollout of the two policies, and detailed administrative data at the worker account level to obtain difference-in-differences estimates of the causal impacts of these interventions.

We distinguish between effects during the access-only period, during the campaign, and afterwards. We can identify the marginal effect of the expansion of access channels, and the marginal effect of the bundle (access and persuasive message). However, we cannot isolate the marginal effect of the ad campaign alone. While this is an important drawback, our objective here is to quantify the effect of lowering transaction costs and the effect of the bundled policies that address both transaction costs and psychological barriers.

Our results show that expanding access alone has small and sometimes insignificant impacts on voluntary contributions. However, we find strong effects for the bundled policies. Relative to municipalities without 7-Eleven presence, those with 7-Eleven access experienced a significant 12% increase in the number of workers making at least one voluntary contribution, and a 10% increase in the total number of voluntary contributions throughout the media campaign. We also find strong evidence of a post-campaign effect for up to seven months, similar in magnitude to the effects during the campaign. This persistence of the bundled policies suggests habit formation for workers making voluntary contributions. These results are maintained across multiple robustness checks.

We then turn to identifying the mechanisms that drive our effects. First, we show that the larger effect of the bundled policies is actually due to the ad campaign, and cannot be simply explained by the access expansion alone via a gradual increase in 7-Eleven usage over time, independent of the campaign. Since the ads were only shown on open, national TV, we

use cable TV penetration to leverage variation in exposure to the campaign, finding much larger results for states with low cable TV penetration (i.e., high ad exposure).

We then consider whether the ads are purely informational by communicating the possibility of making deposits at 7-Eleven. We show, via sample restrictions, that the information channel cannot explain the full effect of the bundled policies, indicating an important role for the persuasion effect. Lastly, we exploit information on deposit channels to present suggestive evidence of the importance of the reminders as the mechanism through which persuasion operates in our context.

Our results provide two key takeaways. First, increasing access matters most if accompanied by a persuasive message. This suggests an important role for bundled policies that address both transaction costs and psychological barriers. Second, we identify strong persistent effects, consistent with habit formation. This suggests that the impact of these policies may be long-lasting, and therefore more cost-effective.

This paper contributes to our understanding of the frictions and biases that affect savings behavior in general, and retirement savings in particular. The previous literature has mostly focused on relaxing one constraint at a time, such as decreasing opening fees and costs (Prina, 2015; Dupas et al., 2018), providing debit cards (Schaner, 2017; Bachas et al., 2018), increasing interest rates (Schaner, 2018; Bertrand et al., 2010), offering commitment devices (Ashraf et al., 2006; Dupas and Robinson, 2013), and providing reminders (Karlan et al., 2016), among others. By focusing on two treatments, we show the importance of potential complementarities between policies, and emphasize that simply addressing transaction costs may not be enough.

We also contribute to our understanding of the long-term impacts on savings behavior. The literature has identified some long-run effects on savings due to, for example, a transitory increase in interest rates (Schaner, 2018), and changes in deposit collection services (De Mel et al., 2013). Additionally, most experimental analyses of persuasive messages have found

that effects are short-lived (DellaVigna and Gentzkow, 2010; Gerber et al., 2011; Simester et al., 2009). We provide robust evidence that persuasion may indeed have long-term effects.

The rest of the paper is organized as follows. Section 2 provides context on the Mexican pension system and the policies analyzed. Section 3 describes the data. Section 4 lays out the empirical strategy. Section 5 presents the results. Section 6 shows evidence disentangling the mechanisms of the effects. Section 7 concludes.

2 Context

2.1 Privatized Social Security in Mexico

In 1997, Mexico transitioned from a pay-as-you-go pension system to a privatized social security system based on individual retirement accounts. The government, through its regulatory agency CONSAR (*Comisión Nacional del Sistema de Ahorro para el Retiro*, National Commission of the Retirement Savings System), oversees this system. Workers choose a manager from a small set of government-approved private fund managers, called AFORES (*Administradoras de Fondos para el Retiro*, Retirement Funds Managers), which generally offer a homogeneous financial product (Duarte and Hastings, 2012; Hastings et al., 2017).⁵ Many well-known banks, insurance companies, and financial institutions have entered the market for AFORES.

The individual retirement fund is made up of three separate accounts. The first, which we call the main account, corresponds to mandatory contributions. Every two months, 6.5% of a worker’s base salary is contributed to this account.⁶ The second account refers to a housing fund, which is administered by INFONAVIT, the Mexican federal institute for

⁵Even though AFORES offer homogeneous products, heterogeneity in labor supply decisions leads to considerable variation in terms of which may be the optimal investment decision. Duarte and Hastings (2012) and Hastings et al. (2017) characterize this heterogeneity and the complexity involved in each worker’s decision-making process.

⁶Workers contribute 1.125%, employers 5.15%, and the federal government 0.225%. The federal government additionally contributes 5.5% of the general minimum wage to each worker’s main account.

workers' housing. Only employers make contributions to the housing fund, at 5% of the base salary on a bimonthly basis.⁷ The third account constitutes voluntary savings. These are any additional contributions made by workers. Voluntary contributions are tax deductible, and may be accessed anytime two months after making the contribution.

Fund managers' investments are constrained by regulations that limit the riskiness of AFORE portfolios. Furthermore, regulations require less risky financial products for portfolios of workers that are nearing retirement age. Notwithstanding these constraints, return rates vary considerably by AFORE. In 2014, AFORE returns ranged from 6 to 12% for the riskier portfolio automatically assigned to workers ages 36 and under, and from 5 to 9% for the least risky financial products automatically assigned to workers over age 59.⁸

In general, access to financial services in Mexico is limited, and returns are fairly low. Only 36% of the adult population has a savings, payroll or investment account in a formal banking institution, and most savings products offer returns below inflation (Peña et al., 2014). In 2014, banks offered an average yearly rate of 2.35%, and up to 3.75% for larger amounts of at least 5 to 10 million pesos (370 to 740 thousand USD).⁹ Note that inflation in Mexico was 4% in 2014. Given the limited options for savings, voluntary contributions are the most attractive choice for most workers in terms of access and returns.

2.2 Policies Incentivizing Voluntary Contributions

Low contribution rates, low base salaries, and extended absences from formal employment present important challenges in this privatized system.¹⁰ At current rates, income replacement rates at retirement will be on average around 40%, with many workers well below that (Villagómez, 2014). In an effort to address this issue, CONSAR has implemented many

⁷This money may be accessed through mortgage loans from INFONAVIT to buy a house, build a house, or remodel an existing property. If workers do not request this money from INFONAVIT, the total amount is added to the main account at the time of retirement.

⁸See, for example, <https://www.dineroenimagen.com/2014-08-06/41473> (last accessed October 23, 2018).

⁹See, for example, <http://www.elfinanciero.com.mx/archivo/quien-paga-mas-por-el-dinero> (last accessed October 23, 2018).

¹⁰In many contexts, including the US, low worker participation rates and low contributions are an important barrier to adequate savings (Skinner, 2007; Benartzi and Thaler, 2013).

policies, some of which have focused on increasing voluntary contributions. Two particular policies in this area are the focus of this paper.

The first policy focused on increasing access points. Until October 2014, workers had two options for making voluntary contributions: they could ask their employer to automatically deduct a fraction directly from their payroll, or they could contact their AFORE representative either in person (for example, at a bank branch) or over the phone.¹¹ Starting on October 6, 2014, CONSAR made voluntary savings available at all 7-Eleven convenience stores. Workers could now make contributions with as little as 50 pesos (3.70 USD) by just providing their national id number to the cashier.¹² The store charges a fixed fee of 4 pesos (0.30 USD) per deposit. With over 1,800 locations, 7-Eleven is an important player in the convenience store market in Mexico, although they lack market presence in 19 out of 32 states.

The second policy, tied to this increase in access, was a national media campaign that ran from July to December 2015. This non-branded advertising consisted most notably of radio and TV ads with a catchy jingle, urging workers to increase their voluntary savings and depicting the possibility of doing this at 7-Eleven locations. Importantly, the ads did not provide any financial information about worker accounts or AFORES, nor a rationale for why workers should increase voluntary savings. The main TV ad is available from the Nielsen-Ibope advertisement archive, with ad identification number 228053.¹³ Figure A1 in the online appendix shows still images from the television ad, encouraging savings in 10-peso coin increments and emphasizing access via 7-Eleven. Table A1 provides the full text of the ad’s message.

It should be noted that other convenience stores and institutions also began accepting voluntary contributions after 7-Eleven access began in 2014. In June 2015, access was expanded

¹¹Starting in June 2014, workers could also make electronic transfers, via their bank accounts. However, this is not a sizable channel, as there are only 35 million bank accounts and 5 million Internet users in Mexico, from a total population of 110 million.

¹²At the time, 1 USD=13.5 pesos.

¹³Nielsen-Ibope offers access to a publicly available database of historical television advertisements at <http://youspot.ibopeagb.com.mx/>. A version of the ad with English subtitles is available at <https://www.youtube.com/watch?v=uSdOpwVJy1o> (last accessed April 4, 2018).

to Telecomm, a government agency that operates telegraphic and satellite services and offers basic financial services. In February 2016, another convenience store chain, Circle-K, began accepting voluntary contributions. By August 2016, access had expanded to Bansefi, a public development bank. We discuss the implications for our empirical strategy in Section 4. Additionally, there was another media campaign in March 2016 that only lasted for a month (Nielsen-Ibope ad number 257082). The characteristics of this campaign were similar to the ones mentioned for the main campaign discussed above.

3 Data

We obtain detailed, anonymized account-level data directly from CONSAR covering a span of 43 months, from January 2013 to July 2016. From the universe of 19 million active worker accounts, we obtain information for all accounts with at least one voluntary contribution over this time period.¹⁴ This gives us a total of 201,565 accounts, from which 75,998 had at least one voluntary contribution prior to October 2014, when 7-Eleven access began. From the remaining 18.8 million worker accounts that did not make any voluntary contributions over these 43 months, we obtain a 10% random sample for computational purposes (1,886,907 accounts).

For each account in our dataset, we observe all voluntary contribution transactions, including the date and amount contributed. We also observe individual-level characteristics (namely, gender, date of birth, state and municipality of residence, and year of affiliation), as well as each worker’s balance in each account in March and September of every year (recall from Section 2 that each worker’s retirement account is made up of three separate accounts: the main account, a housing fund, and the voluntary contributions account). Note that the workers’ residence is only available in our data for the last quarter of 2015, which means we cannot observe workers switching locations over time. We use all this information to

¹⁴CONSAR has a total of 54 million accounts. However, over 35 million are inactive because of sustained absences from the labor market due to a null labor supply, workers entering the informal labor market, or international migration.

create a balanced panel of account-months, where we observe the total number of voluntary contributions and the total amount contributed for each worker in a given month.

We eliminate workers for which we identify inconsistent values within account ids of the gender, date of birth, residence, and year of affiliation variables.¹⁵ Overall, this leaves us with 97% of the original sample or 195,811 accounts with at least one voluntary contribution between 2013 and 2016, with 73,091 accounts making at least one voluntary contribution before October 2014. We are also left with 1,882,599 accounts for which we never observe any voluntary contributions.

Table 1 presents summary statistics for these worker accounts, differentiating between accounts with voluntary contributions prior to 7-Eleven access (“early savers”), accounts with voluntary contributions in October 2014 and after (“treatment savers”), and accounts without voluntary contributions in this time period (“never savers”). Over the 43 months covered in the data, the number of months in which the early savers make voluntary contributions, their total number of voluntary contributions, and their total amount contributed is 3.5 times larger than the treatment savers.

We collapse these data to obtain a balanced panel of municipality-months using the workers’ municipality of residence as determined in the last quarter of 2015. We then join this dataset with a roster of 7-Eleven locations obtained from CONSAR, using locations from the pre-treatment period. Figure 1 shows a map detailing the municipalities where 7-Eleven has locations. Out of 2,298 municipalities with at least one active worker account, 7-Eleven is present in 84 municipalities, mostly in the northern and central parts of the country.¹⁶ Note however that 45% of worker accounts (8,271,542 accounts) correspond to workers living in municipalities with 7-Eleven presence. Our empirical strategy will exploit this geographic variation to identify the causal effect of the policies on savings behavior.

¹⁵Our concern is that inconsistent values may imply an incorrect identification of workers. Note however that we keep worker ids in our sample for which these variables consistently take on a missing value.

¹⁶Although there are around 2,500 municipalities in total, some of them do not have any active worker accounts. These are mostly very rural, poor and small municipalities.

Municipalities with 7-Eleven stores are different than those without, particularly because convenience stores tend to locate in urban areas. We compare state-level characteristics by 7-Eleven presence. Table 2 presents these descriptives using data from the 2010 census. Difference in means tests show that there are substantial differences. States with 7-Eleven presence have a larger population, have individuals that are older (fewer individuals ages 24 and under and more individuals between the ages of 25 and 44), more educated (driven by fewer individuals without schooling), earn higher incomes, and have more access to healthcare. Regardless of level differences across states or municipalities, our identifying assumption below will rely on similar trends over time.

4 Empirical Strategy

4.1 Raw Data Trends

Figure 2 plots different outcomes related to voluntary contributions over time. We present municipality-level aggregates for the 84 municipalities with 7-Eleven presence (darker line) versus the remaining municipalities without 7-Eleven (lighter line). We highlight three relevant time periods: (i) October 2014 to June 2015, when transaction costs decreased through increased access to voluntary savings via 7-Eleven, (ii) July to December 2015, when the access expansion was bundled with non-branded advertising with persuasive reminders to save, and (iii) January to July 2016, when the media campaign was no longer in effect. Outcomes are normalized to one in September 2014 (graphs in levels available upon request), and the difference between the treatment and control states is also shown (dashed line).

Figure 2a considers the total number of active worker accounts that made at least one voluntary contribution in a given month. There is a slight difference between these groups of municipalities in levels prior to October 2014, with 7-Eleven municipalities having a few more workers making at least one voluntary contribution in any given month. The plot of the (normalized) difference suggests very similar trends over time for these two groups in the

pre-treatment period. The plot also indicates a change in the trend post-treatment, which is more apparent in the second time period during the media campaign.

Figure 2b shows a similar graph for the total number of voluntary contributions made in each group of municipalities. Once again, the difference in levels prior to the treatment is virtually nonexistent, and we observe a sharper increase in this outcome for municipalities with 7-Eleven presence once the media campaign begins. Lastly, Figure 2c shows the total amount saved through voluntary contributions. This graph is noisier, and shows a less obvious treatment effect.

Overall, Figure 2 suggests that there was an important change in voluntary contributions in the treatment municipalities relative to the control over the relevant time period, with a stronger difference once the bundled policies were implemented. These plots motivate the empirical strategy, where we estimate the causal effect of the policies by comparing changes over time in municipalities with and without 7-Eleven access.

4.2 Identification Strategy

Our main empirical strategy follows a difference-in-differences (DD) specification at the municipality level. For a balanced panel of municipality-months, we compare changes over time for relevant outcomes in treatment municipalities relative to changes over time in control municipalities. Formally, we estimate the following equation:

$$\begin{aligned}
y_{mt} = & \beta_1(\mathbb{1}_{[7\text{-Eleven}]_m} \times \mathbb{1}_{[\text{access only}]_t}) + \beta_2(\mathbb{1}_{[7\text{-Eleven}]_m} \times \mathbb{1}_{[\text{media campaign}]_t}) \\
& + \beta_3(\mathbb{1}_{[7\text{-Eleven}]_m} \times \mathbb{1}_{[\text{post-campaign}]_t}) + \gamma_m + \theta_t + \varepsilon_{mt}^1
\end{aligned} \tag{1}$$

where y_{mt} is an outcome for municipality m at time t ; $\mathbb{1}_{[7\text{-Eleven}]_m}$ is an indicator for whether or not there is 7-Eleven market presence in municipality m ; $\mathbb{1}_{[\text{access only}]_t}$ is an indicator for the access-only period, from October 2014 to June 2015; $\mathbb{1}_{[\text{media campaign}]_t}$ is an indicator for the media campaign, from July to December 2015; $\mathbb{1}_{[\text{post-campaign}]_t}$ is an indicator for the

months after the campaign, from January to July 2016; γ_m are municipality fixed effects; θ_t are month-year fixed effects; and ε_{mt}^1 is the idiosyncratic error term.

Regressions are weighted by the number of active accounts prior to October 2014 in each municipality.¹⁷ Standard errors are clustered at the municipality level to allow for serial correlation in the unobserved component within municipalities, which is the level at which the treatment varies.

We focus on three main outcomes: (1) the total number of worker accounts with at least one voluntary contribution in a given municipality-month, (2) the total number of voluntary contributions in a municipality-month, and (3) the total amount contributed through voluntary savings in a municipality-month. All outcomes are measured in natural logs to improve model fit and facilitate interpretation.¹⁸ Note that these municipality-level aggregates need only consider the 195,811 accounts with at least one voluntary contribution throughout our sample period. Estimates are identical to specifications that use municipality averages of these outcomes instead, including the accounts that had no voluntary contributions.

The municipality fixed effects in equation 1 imply that we effectively identify coefficients from variation within municipalities over time. This addresses any time-invariant differences between municipalities with and without 7-Eleven stores. The month-year fixed effects address any common trends in savings behavior over time, including yearly seasonality in income and savings behavior.

The DD estimate for the access channel policy alone is given by β_1 . The estimate for the bundled policies (access channel and media campaign) is given by β_2 . The estimate for long-run effects of the bundled policies, specifically once the ad campaign is removed, is represented by β_3 . Under certain conditions, which we discuss below, these DD estimators

¹⁷We classify an account as being active prior to October 2014 if there is at least one non-zero balance in the main account in March 2013, September 2013, March 2014, and September 2014. This excludes accounts that were created or activated post-treatment. We use sampling weights to calculate the total number of accounts from the sample that never made voluntary contributions over this time period.

¹⁸In practice, in order to deal with the potential issue of municipality-months that did not have any voluntary savings, we transform each variable x using the function $\ln(x + 1)$. Results are similar under alternative specifications.

provide the causal effect of increased access to savings channels through 7-Eleven and its interaction with the media campaign.

We are interested not only in the magnitudes of these coefficients, but also in their relative sizes. For each regression, we present two tests of coefficients. First, we test whether the effect of the access channel policy and that of the bundled policies is the same ($H_0 : \beta_1 = \beta_2$). Second, we test whether the effect during and after the campaign is of similar size ($H_0 : \beta_2 = \beta_3$) in order to explore persistence of the effects.

Threats to Identification

The main identifying assumption given our strategy is that outcomes in municipalities without 7-Eleven are a good counterfactual of what would have occurred in municipalities with 7-Eleven access in the absence of the treatments. We argue that workers' unobserved preferences and biases regarding savings behavior are orthogonal to the presence of 7-Eleven stores in their municipality of residence. If these unobservables are fixed over time, they will be completely captured by the municipality fixed effects.

As such, the fundamental source of omitted variable bias are time-varying factors at the municipality level. Therefore, causality depends on assuming that unobservable trends in treatment and control municipalities are similar. If municipalities with 7-Eleven presence are fundamentally different, then this might not hold. To address the concern that the estimated outcomes simply reflect differential trends in treatment and control municipalities, an additional specification also includes treatment-specific quadratic trends of the form $\sum_{n=1}^2 \omega_n (\mathbb{1}_{[7\text{-Eleven}]_m} \times t^n)$ in equation 1.

In order to provide supporting evidence for the parallel pre-trends assumption, and to fully model the treatment effects over time, we also estimate the following equation:

$$y_{mt} = \sum_{k=1}^T \beta_k (\mathbb{1}_{[7\text{-Eleven}]_m} \times \mathbb{1}_{[t=k]}) + \gamma_m + \theta_t + \varepsilon_{mt}^2 \quad (2)$$

where $\mathbb{1}_{[t=k]}$ is an indicator for time period k , T represents the total number of time periods in the regression, and everything else is as defined above. If the parallel pre-trends assumption holds, then the set of estimates $\{\beta_k\}_{k=1}^{\tilde{T}}$, where \tilde{T} represents the last time period prior to implementation of the 7-Eleven access treatment, should all be small and statistically indistinguishable from zero.

A particular threat to identification is that other municipality-level policies may have been introduced during this time period. Recall from Section 2 that three other firms and institutions began accepting voluntary contributions for retirement accounts during this time. Telecomm and Bansefi, both government agencies, entered the voluntary contributions market in June 2015 and August 2016 (outside our sample period), respectively.¹⁹ We do not have information on their locations at the municipality level, but both agencies have locations in all 32 states.²⁰

The other chain of convenience stores that began taking voluntary contributions in February 2016 was Circle-K. This chain is present in 22 states (11 of which coincide with 7-Eleven), although their total number of locations is much smaller than that for 7-Eleven.²¹ Although we do not have municipality-level data for Circle-K, it is likely that their presence at the municipality level overlaps with that of 7-Eleven for the states in which they share locations. Our main concern is therefore that the estimates for the post-campaign period may be confounded with the entry of Circle-K.

To address these issues, we refer to the information on the method and channel used for each transaction. These data indicate that out of all the voluntary contributions made through these new access channels up to July 2016, 82.8% correspond to 7-Eleven, followed by 14.5% at Telecomm, 1.4% at Bansefi, and 1.3% at Circle-K.²² Since voluntary contributions

¹⁹Even though August 2016 is not in our transactions data, anticipatory behavior may bias the estimates.

²⁰Note however that Telecomm presence is more widespread, with an average of 51.5 locations per state, versus 13.5 for Bansefi. Out of 32 states, 18 have fewer than 10 Bansefi locations total. Figure A2 in the online appendix shows the state-level relationship between 7-Eleven presence and Telecomm and Bansefi, showing little correlation.

²¹Circle-K has 1,153 locations, while 7-Eleven has 1,854. Figure A2 in the online appendix shows that there is no correlation between 7-Eleven and Circle-K presence by state.

²²Figure A3 in the online appendix provides relevant histograms for each.

at Circle-K are a very small fraction, we assume that there is no confounding effect, such that the estimates of β_3 are not contaminated by this additional policy.²³ The dynamic DD using equation 2 also lends support to this interpretation.

Furthermore, to the extent that other policies (such as the introduction of the Telecomm and Bansefi channels) were implemented at the national level, they are not a major source of concern for this particular strategy. The prevalence of 7-Eleven as the main source of voluntary contributions from new channels, as well as the campaign that focused widely on 7-Eleven suggest that our empirical strategy is indeed capturing the causal effect of 7-Eleven access on voluntary savings behavior. A robustness check that uses different measures of treatment exposure constructed from the number of 7-Eleven locations also supports this claim.

Another important assumption for our empirical strategy is that there is no differential sorting of workers due to the policy. It seems unlikely that workers would respond by moving to a different municipality given this particular treatment. Note that our only information regarding workers' place of residence corresponds to the last quarter of 2015, as this had not been previously recorded. Our concern would be that some workers responded to the treatment by moving, leading to an incorrect classification of treatment and control workers, and a selected sample for our treatment estimates.

One way to address this is to focus on workers' place of birth instead of residence. We obtain data on state of birth for all workers in our data. Unfortunately, municipality of residence is not recorded. We then use this information to contrast estimates from regressions that classify workers into treatment and control states based on their state of residence versus their state of birth. This allows us to check for any significant differences in the estimates, which would indicate a bias in the residence-based results.

Another related issue would be workers differentially changing their labor supply in response to the policies. This too would lead to a selected sample and possibly an overestimate

²³Note that even if we restrict to Circle-K access dates only (February 2016 and after), the fraction of voluntary contributions there is only 2.4% relative to 7-Eleven's 77.2%. See Figure A3 in the online appendix.

of the effects. To address this we use national survey data collected on a quarterly basis, and show that labor supply outcomes do not vary over time between treatment and control municipalities. These results are available in Figure B1 in the online appendix.

Likewise, we must also rule out the possibility that 7-Eleven expanded differentially in response to the treatment. To the extent that these voluntary contributions do not represent a significant source of revenue for 7-Eleven, it seems unlikely that the firm’s business strategy would be influenced by this policy. Nevertheless, we use 7-Eleven locations from 2014, before the treatment. We were also able to obtain state-level data on locations by the end of 2016. These data indicate that 7-Eleven did not expand into states where they previously had no stores, and that the number of locations per state did not change much.²⁴

5 Effects on Voluntary Contributions

This section presents our findings. We begin by showing the main results using our municipality-level data. We then show some robustness checks: a dynamic DD, a comparison at the state level between assigning treatment based on state of residence versus state of birth, and municipality-level estimates using different variables measuring intensity of treatment.

5.1 Main Results

Table 3 shows the main results from estimating equation 1. We present estimates for the three main outcomes we analyze: (1) the total number of accounts with at least one voluntary contribution in a municipality-month, which we call voluntary savers, (2) the total number of voluntary contributions in a municipality-month, and (3) the total amount contributed voluntarily in a municipality-month, all measured in logs. Estimates are weighted by the number of active worker accounts prior to October 2014 in each municipality, and standard errors are clustered at the municipality level (2,298 municipalities).

²⁴See Figure A4 in the online appendix for more details on changes in 7-Eleven market presence over time.

For each of the outcomes, we start by presenting estimates that include an indicator for 7-Eleven presence instead of municipality fixed effects in columns 1, 5, and 9. Columns 2, 6, and 10 then present the main specification with the municipality fixed effect. Columns 3, 7, and 11 also include treatment-specific quadratic trends. Lastly, columns 4, 8, and 12 remove these quadratic trends and also exclude the Mexico City area.

Our main specification in Table 3 column 2 indicates that relative to baseline there was a significant 5% increase in the number of voluntary savers in municipalities with 7-Eleven market presence during the access-only period (prior to the media campaign), relative to municipalities without 7-Eleven. This was then followed by a significant 12% increase, with respect to baseline levels, during the campaign for municipalities with 7-Eleven access, relative to those without 7-Eleven. For the months after the media campaign, we find a significant 15% differential increase in the number of voluntary savers in treatment municipalities relative to our baseline. A test of coefficients allows us to reject that the access-only (β_1) and the media campaign (β_2) effects are equal. We also reject that the effects during the campaign and after the campaign (β_3) are of the same magnitude.

The main specification for the number of voluntary contributions in column 6 shows a positive, insignificant coefficient for the access-only period, a significant 10% increase during the campaign period, and a significant 12% increase post-campaign. We cannot reject that the magnitude of the effects during and after the campaign is the same. For our third outcome, amount contributed voluntarily, column 10 shows a significant 8% increase during the media campaign, with smaller and insignificant magnitudes for the access-only and post-campaign periods. However, we cannot reject that all three estimates are jointly zero.

As stated above, our key identifying assumption is that municipalities without 7-Eleven presence are a good counterfactual of what would have happened in 7-Eleven municipalities absent the treatment. Therefore, an important concern is differential pre-trends between our treatment and control municipalities. We provide additional evidence below that this does not seem to be the case, by estimating a fully dynamic DD. However, columns 3, 7 and 11

in Table 3 also address this potential issue by including a quadratic trend that is specific to municipalities with 7-Eleven stores.

For the number of voluntary savers in Table 3 column 3, we now find a smaller and statistically insignificant effect for the access-only period. We then find a significant 7% increase during the media campaign in municipalities with 7-Eleven presence relative to those without, and a significant 5% increase post-campaign. We cannot reject that these last two effects are of the same magnitude. For the number of voluntary contributions, we only find a significant 7% differential increase in treatment municipalities relative to the baseline during the ad campaign. Lastly, although the magnitudes in column 11 are larger than those in column 10, there are no significant estimates for the total amount contributed voluntarily and we cannot reject that they are all jointly equal to zero.

Another important concern is whether the Mexico City metropolitan area is the main driver of our results, due to the relatively high presence of 7-Eleven stores and to the types of workers that a large city attracts. The estimates in columns 4, 8 and 12 reestimate the main specification excluding the Mexico City area (we exclude both the entire Federal District and state of Mexico). These results mostly mirror those of the main specification, with large, significant increases in the number of voluntary savers in municipalities with 7-Eleven market presence relative to municipalities without 7-Eleven, and somewhat smaller increases in the number of voluntary contributions. However, for this subsample we cannot reject that the magnitudes of the effects across all outcomes before, during and after the campaign are different from one another.

Taken together, these results suggest three things. First, the initial lowering of transaction costs by increasing access via 7-Eleven had a significant effect on the number of voluntary savers, but not on the number of voluntary contributions nor the amount contributed. Furthermore, this effect is relatively small and loses statistical significance when including treatment-specific trends. In contrast, the effect of the bundled policies during the ad campaign have large, positive and robust effects.

Second, we find strong evidence of persistent effects after the campaign in municipalities with increased access channels. In most specifications, we cannot reject that the size of the effect during the campaign is the same as the effect in the months after the campaign. This suggests an important habit formation in savings behavior that lasts at least up to seven months after bundling with the non-branded advertising campaign.

And third, we find that throughout the implementation of these policies the total amount contributed voluntarily remains unchanged. Additional regressions using the average amount contributed per voluntary saver, the average amount per contribution, and the average number of contributions per saver as the outcome variables show some decline in these measures, especially during and after the media campaign (available upon request). This suggests that although 7-Eleven access and the bundling with the campaign induced an increase in the number of workers making voluntary contributions, contributions became smaller in terms of the amount contributed, thus leading to insignificant effects for that particular outcome.

We present additional results in the online appendix. First, we stratify the sample by workers with and without voluntary contributions prior to the policies. We find effects along both the intensive and extensive margins. We also explore some relevant heterogeneity based on workers' observable characteristics. We show that the results are stronger for men, younger workers, workers that affiliated with social security after 1997, workers with lower retirement account balances, and workers affiliated with fund managers with lower pre-treatment voluntary savings rates.

5.2 Robustness Checks

Dynamic DD

In order to lend support to the parallel pre-trends assumption and to better grasp the dynamics of the effects over time, we estimate equation 2 and present the results graphically in Figure 3. For these regressions, we include municipality and month-year fixed effects as before. These estimates are also weighted by the number of active accounts in each

municipality prior to October 2014. The plots show the β_k coefficients for the interaction of 7-Eleven presence with an indicator for month k , taking September 2014 as the excluded category. Bars show 95% confidence intervals from cluster-robust standard errors.

Figure 3a presents the results for the total number of accounts with voluntary contributions in a given municipality-month. It shows small and (mostly) statistically insignificant coefficients for months prior to October 2014. We cannot reject that all of these estimates are jointly zero. This indicates that there is no substantial difference in the trends of the number of voluntary savers in treatment and control municipalities prior to the policy. This provides reassurance that our key identifying assumption holds.

For the access-only period, all estimates are positive and most are significantly different from zero. Note also that there is some evidence of an increasing trend, which may suggest increased use of the 7-Eleven channel as information is diffused, experience is gained, or trust is built. We further explore this in section 6. On average, we see a 9% differential increase in voluntary savers throughout this period.

During the media campaign, all estimates are positive, significant, and larger than the average of the previous estimates, showing an effect of around 16%. The evidence of an increasing trend is weaker here, with similar-sized coefficients for three of the six months of this period. Lastly, in the post-campaign period, estimates are significant and similar in magnitude to those during the campaign.

Figure 3b shows a similar pattern for the total number of voluntary contributions. Pre-treatment estimates are small, almost all are insignificant, and are jointly indistinguishable from zero. The access-only period shows a few significant estimates, with a bit of evidence of an increasing trend during the last two months of this period. Throughout the media campaign, estimates are larger and are all significantly different from zero. In the post-campaign period, the point estimates are similar in size and significant. Figure 3c shows the estimates for the total amount contributed. These are noisier, are mostly indistinguishable from zero, and lack a clear pattern over time, consistent with our previous findings.

Overall, the estimates in Figure 3 confirm that the relevant outcomes followed a similar trend over time between treatment and control municipalities prior to the implementation of the policies. This helps to validate our identification strategy. They also show that the results in Table 3 are not driven by a handful of significant coefficients. Lastly, the pattern and significance of the point estimates for the post-campaign period lend credibility to the persistence of the effects.

Assigning Treatment Status Based on Place of Birth

Our main results use workers' municipality of residence (as measured in the last quarter of 2015) to assign treatment status based on 7-Eleven market presence. Although unlikely, workers could respond to the treatment by sorting selectively in and out of treatment municipalities. In order to address this, we obtained data on each workers' place of birth. However, these data only specify the state in which they were born, limiting this analysis to the state level.

Figure A5 in the online appendix shows histograms for state of residence and state of birth. Although the differences in density across states between residence and birth are less stark, the overall ranking of states is mostly unchanged. Appendix Figure A6 further explores the relationship between workers' state of residence and their state of birth. For each state of birth, we plot the histogram of state of residence. Except for two states, over 50% of workers born in a given state are residing in that state in 2015, and for many states, this proportion is over 75%.

Table 4 presents estimates of equation 1 using state-level data, comparing between treatment assigned by state of residence (as before) and by state of birth.²⁵ Assuming positively selected migrants, estimates using residence will tend to overstate the effect of the treatments. For all three outcomes, Table 4 shows that results are very similar when using state

²⁵Note that standard errors are clustered by state. However, as there are only 32 states in Mexico, cluster-robust standard errors may be downward-biased. Therefore, we also present wild cluster bootstrap p-values over 1,000 replications to correct for the small number of clusters, following Cameron and Miller (2015).

of residence relative to using state of birth.²⁶ The point estimates are generally smaller when considering state of birth, although we cannot reject that the magnitudes are the same. The fact that these estimates are indistinguishable from one another is reassuring and suggests that our residence-based results do not suffer from significant bias due to selective sorting.

Intensity of Treatment

We now turn to exploiting variation in the intensity of exposure to 7-Eleven market presence. For this, we focus on four measures. First, we use the indicator for whether a municipality has 7-Eleven locations, as before. Second, we use the number of locations by municipality at baseline. Third, we use the number of 7-Eleven stores in a municipality per 100,000 active worker accounts. Lastly, using data from the 2014 economic census, we calculate the 7-Eleven market share by obtaining the fraction of all convenience stores in a municipality that correspond to 7-Eleven.

We use these measures to estimate regressions analogous to equation 1. Table 5 shows the results. The first four columns correspond to the effect on the number of voluntary savers. Except for the number of stores per worker measure, all estimates are statistically significant. The average magnitudes for the effect are smaller for the new intensity measures. For example, the original estimate indicated an average increase of 12% in the number of voluntary savers during the media campaign in municipalities with 7-Eleven presence relative to those without, while these estimates show an average increase of around 1-7%.²⁷

Columns 5 through 8 in Table 5 show the results for the total number of voluntary contributions. Estimates using the intensity measures are generally less significant than the simple indicator. The final four columns present the results for the total amount contributed

²⁶Table A2 in the online appendix shows the main results similar to Table 3 using state of birth to classify workers into treatment and control states. Figure A7 in the online appendix also shows the dynamic DD for this state-level classification.

²⁷We calculate this by taking the estimate and multiplying by the average of the intensity measure. For the full distribution of these measures over the 2,298 municipalities, see Figure A8 in the online appendix.

through voluntary savings. All estimates are statistically indistinguishable from zero, consistent with the findings in Table 3.

The results in Table 5 exploit differing variation in intensity of exposure to the 7-Eleven treatment. Finding a similar pattern across specifications provides reassurance that we are indeed capturing the causal effect of this policy, and not simply a spurious effect from the municipalities with 7-Eleven presence. We cannot determine ex-ante which of these measures is the most relevant in terms of exposure to the treatment, and as such we do not ascribe to a single measure. Note that these estimates provide a clearer picture of the distribution of the effects across space in terms of workers' access to 7-Eleven stores.

6 Disentangling the Mechanisms

Our main results in section 5 indicate a larger effect for the bundled policies than the access expansion on its own. We now turn to disentangling the mechanisms through which these bundled policies may be having such an effect.

Consider the canonical life-cycle earnings model for savings. We present a very simplified toy version to fix some intuition. Workers live in a two-period world, earning income in the first period. They are allowed to transfer income to the future through savings. They maximize lifetime utility subject to their budget constraints in each period:

$$\begin{aligned} \max_{c_1, c_2} \quad & u(c_1) + \beta u(c_2) \\ \text{s.t.} \quad & c_1 + s = y \\ & c_2 = (1 + r)s \end{aligned}$$

where c denotes consumption, β is the time discount factor, s are savings, y is income, and r is the interest rate. The first order condition is then:

$$\beta(1 + r) = \frac{u'(y - s)}{u'((1 + r)s)}$$

This simple model can be easily modified to account for costly savings. Suppose now that individuals choose a savings amount $S = s + k$ in period 1, of which they only receive interest on s and lose k as the transaction cost of saving. Under the modified first order condition, and using the implicit function theorem, it can be shown that the optimal savings rate S^* depends negatively on k and positively on β :

$$\frac{\partial S^*}{\partial k} < 0, \quad \frac{\partial S^*}{\partial \beta} > 0$$

Intuitively, this simply means that if savings are more costly, individuals choose to save less, and if individuals are less present-biased, then they choose to save more.

The interventions we study can be thought of as decreasing k and increasing β . By allowing deposits in a convenient way at 7-Eleven stores, the fixed cost of saving should decline. The uninformative ads of the media campaign serve as reminders that may increase the weight that individuals put on their future in their savings decisions.

Empirically, we face a particular challenge due to the bundling of the policies, as stated above. We can claim an exogenous shift in k through the 7-Eleven access expansion and have the necessary time variation in treatment and control groups. However, the media campaign occurred at the national level, which leaves us without an obvious control group. Under the ideal experiment, the ads would have only been shown in some municipalities, with and without 7-Eleven presence.

Our main results above have identified both $\frac{\partial S^*}{\partial k}$ and $\frac{\partial S^*}{\partial \text{bundle}}$, where the bundle implies shifts in both k and β . Since we cannot identify the marginal effect of the media campaign alone, we focus on contrasting the stand-alone access expansion policy with the bundled policies. Our main findings show that $\frac{\partial S^*}{\partial \text{bundle}} > \frac{\partial S^*}{\partial k}$.

We now turn to disentangling the mechanisms behind this result. First, we discard the possibility that the effect of the bundled policies is simply due to slow and gradual increments in 7-Eleven usage, due to information diffusion, experience and trust. Second, we show that

the media campaign effect is not due to purely an information channel whereby individuals learn about the possibility of 7-Eleven deposits. Lastly, we show suggestive evidence that the reminders are a key element for the effect of the bundled policies.

6.1 Bundled Policies Mechanism

Our main results show a weak 5% increase in the number of voluntary savings during the access-only period, followed by a robust 12% increase during the ad campaign. One interpretation of these results is that the bundled policies have a much stronger effect on voluntary contributions than the access expansion policy. This would suggest that workers are constrained by both transaction costs and psychological biases that limit their savings. However, an alternative interpretation is that individuals learned slowly about the possibility of making deposits at 7-Eleven, and needed time to experience and learn to trust this channel, thus implying that the bundle itself did not really matter. The dynamic DD of the main results suggests that this may be the case.

In order to distinguish between these stories, we exploit state-level data from the 2014 Module on Availability and Use of Information Technologies at Home (MODUTIH). This nationally-representative survey registers each household's number of television sets and cable TV accounts. We construct a measure of cable TV penetration given by the share of televisions with cable in a state, and classify states as high and low exposure to cable TV using the median.²⁸

Since government ads in Mexico can only be shown on open, national television channels (i.e., channels that do not require a cable provider in order to watch them), we assume that our measure of cable TV penetration is monotonically related to exposure to the ads. States with low cable TV usage will thus have experienced a higher exposure to the ad campaign. We stratify our sample by exposure to the ads and estimate equation 1 as before.

²⁸The MODUTIH survey does not register the respondents' municipality of residence, which is why we restrict our analysis to the state level.

Table 6 shows the results for each of the three outcomes. The first four columns focus on the total number of voluntary savers. Columns 1 and 2 correspond to states with low exposure to the ads (high cable TV penetration), while columns 3 and 4 correspond to high ad exposure (low cable TV penetration). Even-numbered columns include average monthly income and hours worked as controls from the quarterly National Occupation and Employment Survey (ENOE). Our preferred specification includes these controls since cable TV penetration may be related to both income and leisure.

Prior the campaign, during the access-only period, point estimates are similar between the high and low exposure states. However, during the media campaign, we find a 5% statistically insignificant increase in low exposure states and a significant 12% increase in high exposure states. This suggests that the increase we observe in the number of voluntary savers during the campaign is indeed due to exposure to the ads, and is not simply reflecting a gradual increase in 7-Eleven usage over time.

Although some of these estimates lack precision, we generally find similar coefficients in high and low exposure states during the access-only period, and larger and more significant coefficients in high exposure states during the campaign across all specifications. These findings indicate that our main results are not just about access, and that the larger effect of the bundled policies is indeed due to the ad campaign.

6.2 Message Mechanism

The findings regarding the bundled policies could be due to the ads serving as persuasive reminders that help overcome psychological biases, or could simply reflect information diffusion about the possibility of making deposits at 7-Eleven locations. We begin by showing that providing information cannot be the sole explanation for our results, and then show suggestive evidence of the importance of the reminders in the ads.

Table 7 provides evidence that this information effect cannot on its own account for the estimates we observe throughout the campaign period. We stratify workers that made volun-

tary contributions during our whole sample period into those with and without transactions during the access-only period, regardless of behavior pre-treatment. We show results from estimating equation 1 based on municipality-level aggregates for the former in columns 1 and 2, and the latter in columns 3 and 4.

The subsample of workers with transactions during the access only period corresponds to the individuals who could *not* have obtained any new information from the ad campaign. Regardless of their savings behavior prior to the treatments, the access channel treatment induced a response on their behalf, such that any subsequent transactions cannot be attributed to them learning about 7-Eleven access.

The estimates in column 1 indicate that the number of workers making at least one voluntary contribution in a given month increased by 5% relative to the pre-treatment period in municipalities with 7-Eleven presence relative to those without. For these same individuals, we then estimate a 10% increase relative to the pre-treatment period in the number of voluntary savers throughout the media campaign. A test of coefficients shows that these magnitudes are indeed different from one another. This differential increase cannot be attributed to an information effect. We then find a 7% increase relative to baseline after the ad campaign. We can reject that this estimate is different from the effect during the campaign, but not different from the access-only coefficient.

We find similar results for the total number of voluntary contributions. There is an insignificant 1% increase during the access-only period, followed by a significant 7% increase during the campaign. These additional deposits cannot be attributed to information diffusion.

Columns 3 and 4 consider municipality-level aggregates from workers without voluntary contributions during the access only period, regardless of their behavior at baseline. For these individuals, we may be confounding the persuasion mechanism with an information effect. Note that by construction we will get large, negative effects during the access only period for this subsample. In column 3, we find that the media campaign increased the

number of workers with at least one voluntary contribution by 21% in municipalities with 7-Eleven presence relative to those without, and then estimate a 37% increase relative to baseline after the campaign. Column 4 shows similar estimates for the total number of transactions.

Aside from the anecdotal evidence that suggests that the first policy was widely covered in the media, the estimates in Table 7 suggest that the mechanism through which the ads have an effect on savings behavior cannot be attributed only to information diffusion about the new access channel at 7-Eleven stores.²⁹ This implies that we are indeed capturing a persuasion effect from the ad campaign.

Persuasion may take many forms, but we provide some suggestive evidence that a key component in this case were the constant reminders to save. For this, we rely on information about which deposit channel was used for voluntary contributions. The transactions data allows us to distinguish between voluntary contributions at access points (mostly 7-Eleven) and all other methods, which include direct contact with AFORES and contributions established directly with employers.³⁰ We partition the data based on the contribution method, obtain the municipality aggregates, and estimate equation 1.

Table 8 presents the results for each of our three main outcomes. Columns 1, 3 and 5 correspond to contributions at access points, while the rest refer to all other contribution methods. The estimates in odd-numbered columns are unsurprisingly large and significant, since prior to treatment there were no contributions at access points.

The results in columns 2, 4 and 6 mirror the main effects presented in Table 3. First, the small and insignificant coefficients for the access-only period serve as a placebo check on other contribution methods. Prior to the campaign, there should be no reason one would expect

²⁹See for example <http://archivo.eluniversal.com.mx/finanzas-cartera/2014/facilitan-aportes-a-la-afore-en-tiendas-7-Eleven-1045548.html> and http://www.milenio.com/negocios/Afore-Consar-Amafore-7Eleven-ahorro-retiro-aportaciones.0_385761554.html (last accessed March 22, 2018).

³⁰We actually only observe whether the contribution occurred at an access point, and the particular channel. Therefore, we cannot distinguish between all other contribution methods.

differences between treatment and control municipalities in transactions that are unrelated to 7-Eleven.

We then find a 10% and 7% increase in the number of voluntary savers and total voluntary contributions, respectively, during the media campaign. These spillovers suggest that the reminders in the ad were effective at increasing savings through channels other than 7-Eleven in treatment municipalities relative to controls.³¹ Given that the ads were shown nation-wide, we posit that this differential increase may be due to the ads serving as constant reminders due to the salience of 7-Eleven stores, even if these additional voluntary contributions were made elsewhere.

The estimates in Tables 7 and 8 indicate that the information component of the message cannot fully explain our main results, and suggest that exposure to constant reminders are an important element of the message. This reinforces the persuasive message as the main mechanism through which the bundling of 7-Eleven expansion and ads affect savings.

7 Conclusion

This paper analyzes the staggered introduction of two policies aimed at increasing savings for retirement. The nature of the implementation allows us to compare the magnitudes of the initial policy that reduced transaction costs for savings by increasing access channels with that of the bundled policies where a media campaign was added. We find that increasing access points by allowing workers to make voluntary contributions at 7-Eleven convenience stores has small to insignificant effects on savings unless accompanied by the ad campaign. We also identify strong persistence of these savings effects after the campaign, suggesting habit formation and that messages may have a long-lasting effect.

We then explore the mechanisms through which the bundled policies have a larger effect. First, we discard the possibility that this is simply an access story with gradual increases

³¹Similar to our findings regarding spillovers, the literature has mostly failed to find evidence of crowd-out effects in savings across a variety of interventions (Dupas and Robinson, 2013; Ashraf et al., 2006; Bachas et al., 2018).

in the utilization of the new access channel. We then provide evidence that the information included in the advertising message cannot fully explain the effects that we find, and show suggestive evidence that reminders played a crucial role in persuading workers to increase savings. Overall, this suggests that the success of the bundled policies is due to their effect on lowering transaction costs and psychological biases that hinder savings.

We make two important contributions. First, we reveal the complementarity between policies aimed at decreasing barriers to savings. The prior literature has focused mostly on analyzing these policies in isolation, when in practice tying them together may matter for their effectiveness. Second, we show that the effects of the bundled policies are long-lived, up to seven months after the ads were implemented. Our findings indicate that media interventions may have a larger value in terms of habit formation than previously thought.

References

- Ashraf, N., D. Karlan, and W. Yin (2006). Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines. *The Quarterly Journal of Economics* 121(2), 635–672.
- Bachas, P., P. Gertler, S. Higgins, and E. Seira (2018). How debit cards enable the poor to save more. Technical report, National Bureau of Economic Research.
- Benartzi, S. and R. H. Thaler (1995). Myopic loss aversion and the equity premium puzzle. *The quarterly journal of Economics* 110(1), 73–92.
- Benartzi, S. and R. H. Thaler (2013). Behavioral economics and the retirement savings crisis. *Science* 339(6124), 1152–1153.
- Bertrand, M., D. Karlan, S. Mullainathan, E. Shafir, and J. Zinman (2010). What’s advertising content worth? Evidence from a consumer credit marketing field experiment. *The Quarterly Journal of Economics* 125(1), 263–306.
- Biggs, A. G. and G. R. Springstead (2008). Alternate measures of replacement rates for social security benefits and retirement income. *Social Security Bulletin* 68, 1.
- Cameron, A. C. and D. L. Miller (2015). A practitioner’s guide to cluster-robust inference. *Journal of Human Resources* 50(2), 317–372.
- Carroll, G. D., J. J. Choi, D. Laibson, B. C. Madrian, and A. Metrick (2009). Optimal defaults and active decisions. *The quarterly journal of economics* 124(4), 1639–1674.
- Chen, M. K. (2013). The effect of language on economic behavior: Evidence from savings rates, health behaviors, and retirement assets. *American Economic Review* 103(2), 690–731.

- Chetty, R., J. N. Friedman, S. Leth-Petersen, T. H. Nielsen, and T. Olsen (2014). Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark. *The Quarterly Journal of Economics* 129(3), 1141–1219.
- Choi, J. J., D. Laibson, B. C. Madrian, and A. Metrick (2003). Optimal defaults. *American Economic Review* 93(2), 180–185.
- Choi, J. J., D. Laibson, B. C. Madrian, and A. Metrick (2005). Saving for retirement on the path of least resistance. Technical report, National Bureau of Economic Research.
- De Mel, S., C. McIntosh, and C. Woodruff (2013). Deposit collecting: Unbundling the role of frequency, salience, and habit formation in generating savings. *American Economic Review* 103(3), 387–92.
- DellaVigna, S. and M. Gentzkow (2010). Persuasion: empirical evidence. *Annual Review of Economics* 2(1), 643–669.
- Duarte, F. and J. S. Hastings (2012). Fettered consumers and sophisticated firms: Evidence from Mexico’s privatized social security market. Technical report, National Bureau of Economic Research.
- Duflo, E. and E. Saez (2002). Participation and investment decisions in a retirement plan: The influence of colleagues choices. *Journal of public Economics* 85(1), 121–148.
- Dupas, P., D. Karlan, J. Robinson, and D. Ubfal (2018). Banking the unbanked? Evidence from three countries. *American Economic Journal: Applied Economics* 10(2), 257–97.
- Dupas, P. and J. Robinson (2013). Why don’t the poor save more? Evidence from health savings experiments. *American Economic Review* 103(4), 1138–71.
- Gerber, A. S., J. G. Gimpel, D. P. Green, and D. R. Shaw (2011). How large and long-lasting are the persuasive effects of televised campaign ads? Results from a randomized field experiment. *American Political Science Review* 105(1), 135–150.

- Handel, B. and J. Schwartzstein (2018). Frictions or mental gaps: What’s behind the information we (don’t) use and when do we care? *Journal of Economic Perspectives* 32(1), 155–78.
- Hastings, J., A. Hortaçsu, and C. Syverson (2017). Sales force and competition in financial product markets: The case of Mexico’s social security privatization. *Econometrica* 85(6), 1723–1761.
- Hastings, J. S. and L. Tejeda-Ashton (2008). Financial literacy, information, and demand elasticity: survey and experimental evidence from Mexico. Technical report, National Bureau of Economic Research.
- Illanes, G. (2016). Switching costs in pension plan choice. Technical report, Northwestern University Working Paper.
- Iyengar, S. S., G. Huberman, and W. Jiang (2004). How much choice is too much? Contributions to 401 (k) retirement plans. *Pension design and structure: New lessons from behavioral finance* 83, 84–87.
- Karlan, D., M. McConnell, S. Mullainathan, and J. Zinman (2016). Getting to the top of mind: How reminders increase saving. *Management Science* 62(12), 3393–3411.
- Karlan, D. and J. Morduch (2010). Access to finance. In *Handbook of development economics*, Volume 5, pp. 4703–4784. Elsevier.
- Karlan, D., A. L. Ratan, and J. Zinman (2014). Savings by and for the poor: A research review and agenda. *Review of Income and Wealth* 60(1), 36–78.
- Krasnokutskaya, E. and P. Todd (2009). Investor behavior and fund performance under a privatized retirement accounts system: Evidence from Chile. Technical report, University of Michigan, Michigan Retirement Research Center.

- Looney, C. A. and A. M. Hardin (2009). Decision support for retirement portfolio management: Overcoming myopic loss aversion via technology design. *Management Science* 55(10), 1688–1703.
- Luco, F. (2013). Switching costs and competition in retirement investment. Technical report, The Center for the Study of Industrial Organization-Northwestern University. The Center for the Study of Industrial Organization-Northwestern University.
- Morrissey, M. (2016). The state of American retirement. Technical report, Economic Policy Institute, Washington, DC.
- Pallares-Miralles, M., C. Romero, and E. Whitehouse (2012). International patterns of pension provision II: a worldwide overview of facts and figures.
- Peña, X., C. Hoyo, D. Tuesta, et al. (2014). Determinantes de la inclusión financiera en México a partir de la ENIF 2012. Technical report, BBVA Research.
- Prina, S. (2015). Banking the poor via savings accounts: Evidence from a field experiment. *Journal of Development Economics* 115, 16–31.
- Schaner, S. (2017). The cost of convenience? Transaction costs, bargaining power, and savings account use in Kenya. *Journal of Human Resources* 52(4), 919–945.
- Schaner, S. (2018). The persistent power of behavioral change: Long-run impacts of temporary savings subsidies for the poor. *American Economic Journal: Applied Economics* 10(3), 67–100.
- Simester, D., Y. J. Hu, E. Brynjolfsson, and E. T. Anderson (2009). Dynamics of retail advertising: Evidence from a field experiment. *Economic Inquiry* 47(3), 482–499.
- Skinner, J. (2007). Are you sure you’re saving enough for retirement? *Journal of Economic Perspectives* 21(3), 59–80.

Villagómez, F. A. (2014). El ahorro para el retiro. Una reflexión para México. *El trimestre económico* 81(323), 549–576.

Whitehouse, E. (2006). *Pensions panorama: Retirement-income systems in 53 countries*. The World Bank.

Figures and Tables

Table 1:
Worker account descriptive statistics

	Early savers (1)	Treatment savers (2)	Never savers (3)
Share female	0.37 (0.48)	0.39 (0.49)	0.36 (0.48)
Age in January 2013	45.17 (13.02)	39.47 (12.39)	35.13 (11.37)
Year of affiliation	1990 (12.79)	1995 (12.49)	1998 (10.75)
Account balance March 2013 (MXN):			
Main account	163,709.18 (222,100.11)	172,064.91 (409,014.07)	74,513.47 (112,381.44)
Housing account	49,616.61 (89,752.19)	38,835.67 (71,491.02)	21,806.63 (42,000.08)
Voluntary contributions	28,531.07 (102,709.52)	2,502.49 (14,335.78)	161.29 (1,388.61)
Account balance March 2016 (MXN):			
Main account	185,283.02 (243,788.51)	150,466.68 (207,432.94)	99,058.48 (136,807.24)
Housing account	58,704.96 (108,685.21)	50,121.55 (87,074.52)	29,154.24 (51,249.51)
Voluntary contributions	47,645.53 (156,419.19)	11,108.88 (41,759.21)	574.57 (4,496.99)
Months with voluntary contributions	8.06 (11.63)	2.29 (3.58)	0.00 (0.00)
Total voluntary contributions	13.51 (25.14)	3.67 (7.71)	0.00 (0.00)
Total amount contributed (MXN thousands)	39.97 (102.76)	11.39 (38.86)	0.00 (0.00)
Total accounts	73,091	122,720	1,882,599

Notes: Means are shown with standard deviations in parentheses. Accounts are stratified based on voluntary contributions: accounts with voluntary contributions prior to October 2014 (column 1), those with contributions after 7-Eleven access began (column 2), and those without any contributions (column 3). The transactions data spans 43 months from January 2013 to July 2016.

Table 2:
State-level descriptives by 7-Eleven presence

	Without 7-Eleven	With 7-Eleven
Population, millions	2.74 (1.88)	4.65* (3.92)
Fraction female	0.51 (0.01)	0.51 (0.01)
Fraction ages 0-14	0.30 (0.02)	0.28** (0.02)
Fraction ages 15-24	0.19 (0.01)	0.18** (0.01)
Fraction ages 25-44	0.28 (0.02)	0.30** (0.02)
Fraction ages 45-64	0.15 (0.01)	0.16 (0.02)
Fraction age 65 and above	0.07 (0.01)	0.07 (0.01)
Fraction no schooling	0.10 (0.04)	0.06** (0.03)
Fraction elementary school	0.68 (0.04)	0.67 (0.08)
Fraction high school	0.14 (0.03)	0.16 (0.04)
Fraction college	0.08 (0.03)	0.10 (0.06)
Years of education	7.01 (0.71)	7.52* (0.92)
Income, thousands	3.79 (1,287.61)	4.73* (1,781.49)
Fraction receiving government aid	0.28 (0.09)	0.23 (0.09)
Fraction receiving remittances	0.03 (0.02)	0.02 (0.02)
Fraction with access to healthcare	0.34 (0.11)	0.43* (0.15)
Observations	19	13

Notes: This table shows descriptive statistics from the 2010 census at the state level, distinguishing states by 7-Eleven presence. Averages and standard deviation for both groups of states shown. Stars denote significance for a difference in means test.

*** p<0.01, ** p<0.05, * p<0.1

Table 3:
Effect of treatments on voluntary contributions

	Voluntary savers				Voluntary contributions				Amount (in thousands of pesos)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
7-Eleven × access only	0.0483** (0.022)	0.0483** (0.022)	0.0275 (0.018)	0.0737*** (0.022)	0.0265 (0.020)	0.0265 (0.020)	0.0172 (0.017)	0.0491** (0.020)	0.0337 (0.039)	0.0337 (0.039)	0.0534 (0.050)	0.0753 (0.065)
7-Eleven × media campaign	0.122*** (0.026)	0.122*** (0.027)	0.0680*** (0.024)	0.0996*** (0.032)	0.103*** (0.028)	0.103*** (0.028)	0.0662*** (0.024)	0.0725** (0.031)	0.0757** (0.037)	0.0757** (0.037)	0.0941 (0.075)	0.0911 (0.099)
7-Eleven × post-campaign	0.149*** (0.033)	0.149*** (0.034)	0.0494* (0.028)	0.0809** (0.036)	0.115*** (0.035)	0.115*** (0.036)	0.0391 (0.027)	0.0533* (0.032)	0.0454 (0.043)	0.0454 (0.044)	0.0562 (0.113)	0.00123 (0.147)
Observations	98,814	98,814	98,814	92,751	98,814	98,814	98,814	92,751	98,814	98,814	98,814	92,751
R-squared	0.321	0.987	0.987	0.986	0.311	0.985	0.985	0.984	0.277	0.938	0.938	0.932
Indicator for 7-Eleven Municipality FE	X	X	X	X	X	X	X	X	X	X	X	X
Treatment-specific trends			X				X				X	
Excluding Mexico City				X				X				X
Coefficient tests:												
$H_0 : \beta_1 = \beta_2$	0.00	0.00	0.01	0.14	0.00	0.00	0.01	0.26	0.29	0.30	0.44	0.83
$H_0 : \beta_2 = \beta_3$	0.07	0.08	0.17	0.27	0.50	0.50	0.07	0.31	0.36	0.36	0.50	0.20
Mean dep. variable	111.22	111.22	111.22	66.28	180.12	180.12	180.12	116.11	641.18	641.18	641.18	330.92

Notes: This table shows the main results from expanding 7-Eleven access before, during, and after the media campaign. Observations are at the municipality-month level. Voluntary savers are the total number of accounts in a municipality-month with at least one voluntary contribution. Voluntary contributions are the total number of contributions to voluntary accounts. Amount contributed is the total amount from voluntary contributions (measured in thousands of Mexican pesos). The outcome variables are transformed into logs for the estimation. The mean of the dependent variable for the treatment municipalities at baseline is shown. Regressions include month-year FE, and are weighted by the number of accounts pre-treatment. Robust standard errors clustered by municipality in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 4:
State-level estimates of the effect of treatments on voluntary
contributions: State of residence vs state of birth

	Voluntary savers		Voluntary contributions		Amount (in thousands of pesos)	
	(1)	(2)	(3)	(4)	(5)	(6)
7-Eleven × access only	0.0306 (0.041) [0.460]	0.0346 (0.028) [0.250]	0.0256 (0.036) [0.490]	0.0261 (0.026) [0.340]	0.00363 (0.047) [0.940]	-0.00643 (0.033) [0.870]
7-Eleven × media campaign	0.102** (0.049) [0.040]	0.0875** (0.031) [0.020]	0.0959* (0.051) [0.070]	0.0835** (0.034) [0.030]	0.0667* (0.037) [0.060]	0.0393 (0.044) [0.410]
7-Eleven × post-campaign	0.111* (0.055) [0.050]	0.0729* (0.036) [0.070]	0.102* (0.057) [0.080]	0.0570 (0.039) [0.160]	0.119* (0.061) [0.070]	0.0284 (0.058) [0.610]
Observations	1,376	1,376	1,376	1,376	1,376	1,376
R-squared	0.994	0.997	0.994	0.997	0.973	0.979
State of residence	X		X		X	
State of birth		X		X		X
Mean dep. variable	854.43	816.10	1,383.63	1,320.31	4,745.39	4,609.20

Notes: This table shows state-level results from expanding 7-Eleven access before, during, and after the media campaign. Odd columns classify workers into treatment and control based on their state of residence, even columns use state of birth. Observations are at the state-month level. Voluntary savers are the total number of accounts in a state-month with at least one voluntary contribution. Voluntary contributions are the total number of contributions to voluntary accounts. Amount contributed is the total amount from voluntary contributions (measured in thousands of Mexican pesos). The outcome variables are transformed into logs for the estimation. The mean of the dependent variable for the treatment states at baseline is shown. Regressions include state and month-year FE, and are weighted by the number of accounts pre-treatment. Robust standard errors clustered by state in parentheses. Wild cluster bootstrap p-values in brackets (stars denote significance from these values).

*** p<0.01, ** p<0.05, * p<0.1

Table 5:
Effect of treatments on voluntary contributions: Intensity of
treatment

	Voluntary savers				Voluntary contributions				Amount (in thousands of pesos)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
7-Eleven measure \times access only	0.0483** (0.022)	0.0005*** (0.000)	0.0005 (0.000)	0.0160*** (0.006)	0.0265 (0.020)	0.0003** (0.000)	0.0003 (0.000)	0.0105* (0.006)	0.0337 (0.039)	0.0000 (0.000)	0.0000 (0.000)	0.0034 (0.010)
7-Eleven measure \times media campaign	0.1225*** (0.027)	0.0005** (0.000)	0.0006 (0.000)	0.0207*** (0.007)	0.1031*** (0.028)	0.0003 (0.000)	0.0004 (0.000)	0.0130* (0.007)	0.0757** (0.037)	0.0001 (0.001)	0.0006 (0.001)	0.0149 (0.014)
7-Eleven measure \times post-campaign	0.1485*** (0.034)	0.0005* (0.000)	0.0005 (0.000)	0.0177** (0.009)	0.1147*** (0.036)	0.0002 (0.000)	0.0002 (0.000)	0.0097 (0.010)	0.0454 (0.044)	-0.0001 (0.001)	0.0002 (0.001)	0.0055 (0.020)
Observations	98,814	98,814	98,814	98,814	98,814	98,814	98,814	98,814	98,814	98,814	98,814	98,814
R-squared	0.987	0.987	0.987	0.987	0.985	0.985	0.985	0.985	0.938	0.938	0.938	0.938
Intensity measures:												
Indicator	X				X				X			
Num. stores		X				X				X		
Stores per acct.			X				X				X	
Market share				X				X				X
Coefficient tests:												
$H_0 : \beta_1 = \beta_2$	0.00	0.72	0.33	0.27	0.00	0.55	0.59	0.63	0.30	0.88	0.33	0.44
$H_0 : \beta_2 = \beta_3$	0.08	0.88	0.34	0.35	0.50	0.86	0.33	0.46	0.36	0.52	0.44	0.41
Mean dep. var.	111.22	111.22	111.22	111.22	180.12	180.12	180.12	180.12	641.18	641.18	641.18	641.18
Mean intensity measure	1.00	22.07	91.28	3.19	1.00	22.07	91.28	3.19	1.00	22.07	91.28	3.19

Notes: This table shows the main results from expanding 7-Eleven access before, during, and after the media campaign, using various measures for intensity of 7-Eleven exposure. Observations are at the municipality-month level. Voluntary savers are the total number of accounts in a municipality-month with at least one voluntary contribution. Voluntary contributions are the total number of contributions to voluntary accounts. Amount contributed is the total amount from voluntary contributions (measured in thousands of Mexican pesos). Four measures of intensity are presented: (i) an indicator for 7-Eleven presence, (ii) the total number of stores in the municipality, (iii) the number of stores per 100,000 worker accounts pre-treatment, and (iv) 7-Eleven's market share, defined as the fraction of convenience stores that belong to 7-Eleven. The outcome variables are transformed into logs for the estimation. The mean of the dependent variable for the treatment municipalities at baseline is shown. Regressions include municipality and month-year FE, and are weighted by the number of accounts pre-treatment. Robust standard errors clustered by municipality in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 6:
Effect of treatments on voluntary contributions: Stratifying by
exposure to the media campaign

	Voluntary savers				Voluntary contributions				Amount (in thousands of pesos)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
7-Eleven \times access only	0.0308 (0.063) [0.680]	0.0249 (0.063) [0.730]	0.0226 (0.060) [0.750]	0.0287 (0.057) [0.650]	0.00642 (0.060) [0.920]	0.00198 (0.061) [0.970]	0.0284 (0.052) [0.610]	0.0330 (0.051) [0.560]	0.0936 (0.105) [0.480]	0.0935 (0.105) [0.490]	-0.0211 (0.057) [0.770]	-0.0171 (0.056) [0.800]
7-Eleven \times media campaign	0.0468 (0.094) [0.670]	0.0483 (0.092) [0.650]	0.118 (0.067) [0.110]	0.124* (0.062) [0.090]	0.0431 (0.092) [0.680]	0.0448 (0.091) [0.660]	0.114 (0.070) [0.140]	0.118 (0.069) [0.140]	0.0258 (0.077) [0.750]	0.0283 (0.076) [0.730]	0.107** (0.047) [0.020]	0.114** (0.050) [0.030]
7-Eleven \times post-campaign	0.0145 (0.134) [0.930]	-0.00476 (0.127) [0.970]	0.152** (0.065) [0.040]	0.156** (0.058) [0.030]	0.0136 (0.140) [0.930]	-0.00378 (0.134) [0.980]	0.143* (0.069) [0.070]	0.146* (0.063) [0.060]	0.0824 (0.109) [0.450]	0.0710 (0.106) [0.520]	0.173* (0.071) [0.050]	0.169** (0.070) [0.030]
Observations	688	688	688	688	688	688	688	688	688	688	688	688
R-squared	0.972	0.973	0.995	0.995	0.977	0.978	0.995	0.995	0.871	0.871	0.980	0.980
Controls		X		X		X		X		X		X
Exposure to ads	Low	Low	High	High	Low	Low	High	High	Low	Low	High	High
Mean dependent variable	335.91	335.91	1178.51	1178.51	619.06	619.06	1861.49	1861.49	1382.29	1382.29	6847.33	6847.33

Notes: This table shows the results from expanding 7-Eleven access before, during, and after the media campaign, splitting the data by exposure to the ads. Observations are at the state-month level. Exposure to the ads is measured in terms of cable TV penetration (share of televisions with cable) at the state level in 2014. High exposure corresponds to cable TV penetration below the median. Voluntary savers are the total number of accounts in a state-month with at least one voluntary contribution. Voluntary contributions are the total number of contributions to voluntary accounts. Amount contributed is the total amount from voluntary contributions (in thousands of Mexican pesos). The outcome variables are transformed into logs for the estimation. The mean of the dependent variable for the treatment states at baseline is shown. Regressions include state and month-year FE, and are weighted by the number of accounts pre-treatment. Controls are average income and number of hours worked. Robust standard errors clustered by state in parentheses. Wild cluster bootstrap p-values in brackets (starts denote significance from these values).

*** p<0.01, ** p<0.05, * p<0.1

Table 7:
Effect of treatments on voluntary contributions: Stratifying by
activity during access only period

	Accounts with VCs during access only period		Accounts without VCs during access only period	
	Vol. savs. (1)	Vol. conts. (2)	Vol. savs. (3)	Vol. conts. (4)
7-Eleven \times access only	0.0521** (0.024)	0.0101 (0.021)	-1.665*** (0.196)	-1.774*** (0.210)
7-Eleven \times media campaign	0.0971*** (0.027)	0.0660** (0.028)	0.209*** (0.042)	0.237*** (0.051)
7-Eleven \times post-campaign	0.0685** (0.030)	0.0461 (0.034)	0.372*** (0.048)	0.372*** (0.055)
Observations	98,814	98,814	98,814	98,814
R-squared	0.989	0.986	0.913	0.913
Coefficient tests:				
$H_0 : \beta_1 = \beta_2$	0.00	0.00	0.00	0.00
$H_0 : \beta_2 = \beta_3$	0.06	0.24	0.00	0.00
$H_0 : \beta_1 = \beta_3$	0.43	0.15	0.00	0.00
Mean dep. variable	74.29	127.39	36.92	52.73

Notes: This table shows the results from expanding 7-Eleven access before, during, and after the media campaign, splitting the data by accounts with and without voluntary contributions (VCs) during the access only (pre-campaign) period. Observations are at the municipality-month level. Voluntary savers are the total number of accounts in a municipality-month with at least one voluntary contribution. Voluntary contributions are the total number of contributions to voluntary accounts. The outcome variables are transformed into logs for the estimation. The mean of the dependent variable for the treatment municipalities at baseline is shown. Regressions include municipality and month-year FE, and are weighted by the number of accounts pre-treatment. Robust standard errors clustered by municipality in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

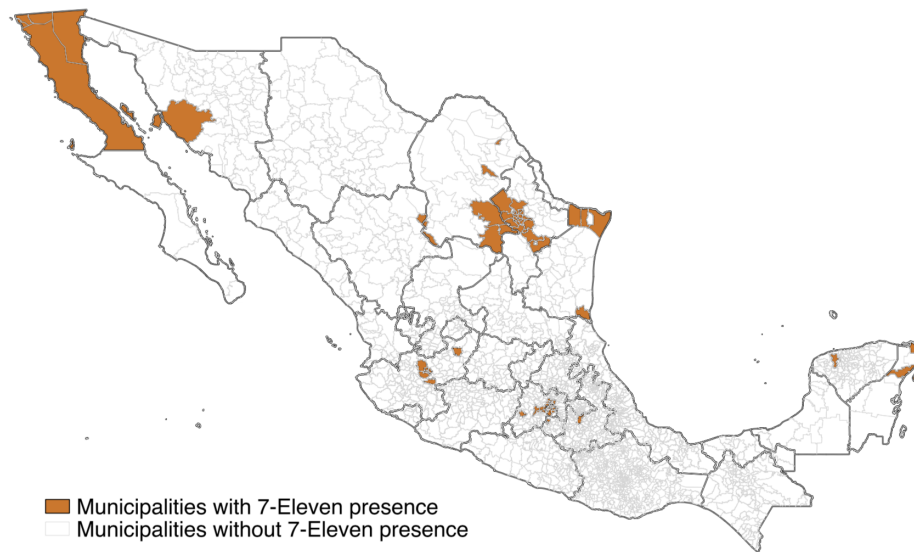
Table 8:
Effect of treatments on voluntary contributions: Access points vs
other contribution methods

	Voluntary savers		Voluntary contributions		Amount (in thousands of pesos)	
	(1)	(2)	(3)	(4)	(5)	(6)
7-Eleven \times access only	1.637*** (0.114)	0.0334 (0.022)	1.824*** (0.128)	0.00746 (0.02)	1.290*** (0.11)	0.0301 (0.039)
7-Eleven \times media campaign	2.003*** (0.123)	0.0999*** (0.027)	2.252*** (0.137)	0.0717** (0.028)	1.742*** (0.134)	0.0671* (0.037)
7-Eleven \times post-campaign	2.175*** (0.136)	0.126*** (0.034)	2.451*** (0.146)	0.0800** (0.036)	2.141*** (0.136)	0.0355 (0.044)
Observations	98,814	98,814	98,814	98,814	98,814	98,814
R-squared	0.849	0.987	0.849	0.985	0.795	0.938
VCs at access points (7-Eleven)	X		X		X	
VCs for all other contribution methods		X		X		X
Coefficient tests:						
$H_0 : \beta_1 = \beta_2$	0.00	0.00	0.00	0.00	0.00	0.36
$H_0 : \beta_2 = \beta_3$	0.00	0.08	0.00	0.63	0.00	0.35
Mean dependent variable	0.00	111.22	0.00	180.12	0.00	641.18

Notes: This table shows the main results from expanding 7-Eleven access before, during, and after the media campaign, splitting the data by the method in which workers make the contribution. Odd columns correspond to voluntary contributions (VCs) at access points (mainly, 7-Eleven). Even columns correspond to all other contribution methods. Observations are at the municipality-month level. Voluntary savers are the total number of accounts in a municipality-month with at least one voluntary contribution. Voluntary contributions are the total number of contributions to voluntary accounts. Amount contributed is the total amount from voluntary contributions (measured in thousands of Mexican pesos). The outcome variables are transformed into logs for the estimation. The mean of the dependent variable for the treatment municipalities at baseline is shown. Regressions include municipality and month-year FE, and are weighted by the number of accounts pre-treatment. Robust standard errors clustered by municipality in parentheses.


*** p<0.01, ** p<0.05, * p<0.1

Figure 1:
7-Eleven presence by municipality




Notes: This map shows the distribution of 7-Eleven presence in Mexico by municipality.

Figure 2:
Voluntary contributions over time by 7-Eleven presence



(a) Accounts with voluntary contributions

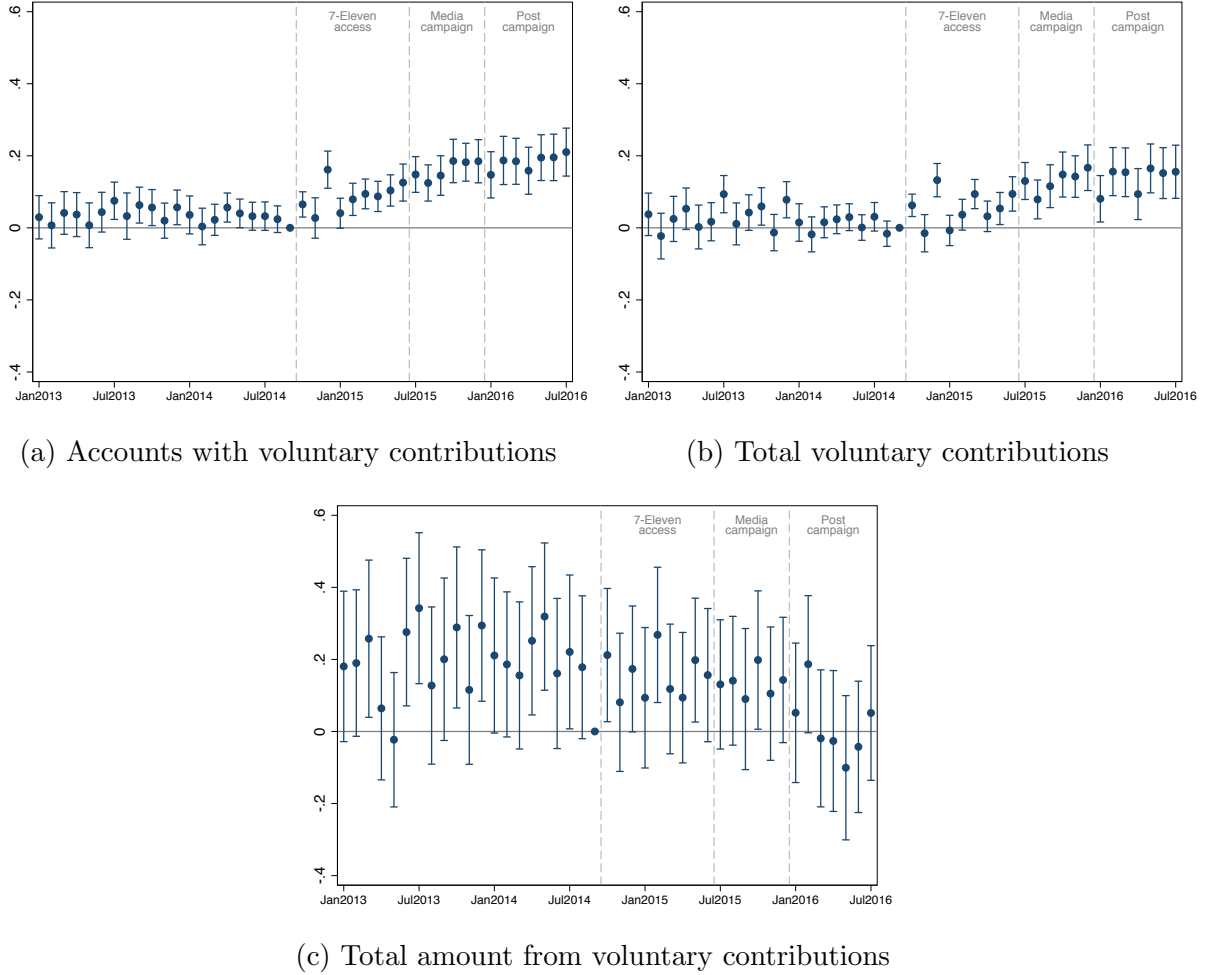
(b) Total voluntary contributions



(c) Total amount from voluntary contributions

Notes: These plots show the raw data, divided by municipalities with and without 7-Eleven presence. The first graph shows the total number of accounts in a municipality with at least one voluntary contribution in a given month, the second graph shows the total number of voluntary contributions, and the third graph shows the total amount contributed to the voluntary savings account. All data series have been normalized so that the value in September 2014 is equal to one. The vertical lines show each of the treatments: 7-Eleven access only, access during the campaign, and access after the campaign.

Figure 3:
Effect of treatments on voluntary contributions



Notes: These plots show the main results from expanding 7-Eleven access before, during, and after the media campaign, within a dynamic DD framework. Observations are at the municipality-month level. The first panel shows the total number of accounts in a municipality-month with at least one voluntary contribution, the second shows the total number of contributions to voluntary accounts, and the third corresponds to the total amount from voluntary contributions (measured in thousands of Mexican pesos). The outcome variables are transformed into logs for the estimation. Coefficients for month indicators interacted with 7-Eleven presence are shown, from regressions that include municipality and month-year FE, and are weighted by the number of accounts pre-treatment. Bars correspond to 95% confidence intervals based on robust standard errors clustered by municipality.

Appendix for Online Publication

A Additional Figures and Tables

We present additional supporting evidence in this appendix. Table A1 displays the text of the media campaign ad in the original language and an English translation. This supports the claim that no financial information was provided in terms of why retirement savings matter, as well as the term of the savings instrument. Figure A1 presents some still images from the advertisement, showing the different elements emphasized by the campaign (i.e., saving in small increments and making deposits at convenience stores).

Figure A2 shows the correlation between 7-Eleven presence and other locations that also began accepting voluntary contributions at the state level. Figure A3 displays histograms of the number of contributions made at 7-Eleven and these other locations over time. Figure A4 shows changes in 7-Eleven locations at the state level over time (2014 versus 2016).

Figure A5 shows the distribution of workers' state of birth and state of residence as described in the data. Figure A6 shows the correlation between workers' state of residence and state of birth.

Table A2 shows the main results from estimating equation 1 as in the main text, at the state level instead of municipality-level aggregates. Figure A7 shows the main estimates under the dynamic framework of equation 2 using state-level aggregates.

Lastly, Figure A8 displays histograms for each of the 7-Eleven intensity measures used in the robustness check presented in the main text.

Table A1:
Ad campaign text

A-2

<u>Original Spanish text</u>	<u>English translation</u>
10 pesitos al día muy fácil tú puedes ahorrar	10 pesos each day is easy to save
Poco a poco un retiro más digno vas a asegurar	And little by little you will ensure a better retirement
10 morlacos, 10 varos o como les quieras llamar	10 bucks, 10 clams, whatever you call them
Es sencillo lograrlo sin tu bolsillo afectar	It's easy to achieve, it's no burden on your pocket
Con 10 pesitos (diez diez) para tu AFORE (diez diez)	With just 10 pesos (ten, ten), for your AFORE (ten, ten)
Lo de hoy es ahorrar y después tu futuro gozar	It's trendy to save, so you can then enjoy your future
Con 10 pesitos (diez diez) para tu AFORE (diez diez)	With just 10 pesos (ten, ten), for your AFORE (ten, ten)
Hay que ahorrar diariamente con 10 pesitos o más	You must save everyday, just 10 pesos or more
Súmale 10 pesitos al día para asegurar	Add 10 pesos a day in order to ensure
Tu futuro, tu AFORE y muy buena pensión alcanzar	that your future, AFORE and pension will turn out right
En Seven Eleven y en Telecomm tú podrás aportar	At 7-Eleven and Telecomm you can save
Deposita sin costo 50 pesitos o más	Deposit without charge from 50 pesos or more
Con 10 pesitos (diez diez) para tu AFORE (diez diez)	With just 10 pesos (ten, ten), for your AFORE (ten, ten)
Lo de hoy es ahorrar y después tu futuro gozar	It's trendy to save, so you can then enjoy your future
Con 10 pesitos (diez diez) para tu AFORE (diez diez)	With just 10 pesos (ten, ten), for your AFORE (ten, ten)
Hay que ahorrar diariamente con 10 pesitos o más	You must save everyday, just 10 pesos or more
Se te nota que sí traes morralla, tú puedes ahorrar	One can see that you've got spare change, you can save
El guardián de tu AFORE y tu aliado sin duda es CONSAR.	Your best ally and guard for your AFORE, without a doubt, is CONSAR.

Notes: The left column shows the original Spanish text from the television ad's jingle. The right column shows the equivalent in English (authors' own translation). The ad uses the diminutive form of the word "peso" to emphasize that workers do not need to make large contributions. Note also that even though AFORE is the acronym for the retirement fund managers, it is customary to refer to one's retirement account as an AFORE as well (for example, "10 pesos for your AFORE" means 10 pesos for your individual retirement account).

Table A2:
Effect of treatments on voluntary contributions: State-level
estimates using state of birth

	Voluntary savers (in thousands)				Voluntary contributions (in thousands)				Amount contributed (in millions of pesos)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
7-Eleven × access only	0.0343 (0.028) [0.260]	0.0343 (0.028) [0.260]	0.0179 (0.026) [0.500]	0.0464 (0.030) [0.130]	0.0259 (0.026) [0.340]	0.0259 (0.026) [0.340]	0.0219 (0.020) [0.280]	0.0368 (0.027) [0.210]	-0.00655 (0.033) [0.870]	-0.00655 (0.033) [0.870]	-0.0332 (0.048) [0.510]	-0.00124 (0.039) [0.970]
7-Eleven × media campaign	0.0870** (0.031) [0.020]	0.0870** (0.031) [0.020]	0.0620** (0.023) [0.010]	0.0806* (0.037) [0.060]	0.0832** (0.034) [0.030]	0.0832** (0.034) [0.030]	0.0769*** (0.023) [0.000]	0.0753* (0.041) [0.090]	0.0392 (0.044) [0.410]	0.0392 (0.044) [0.410]	-0.0289 (0.092) [0.760]	0.0295 (0.048) [0.550]
7-Eleven × post-campaign	0.0721* (0.035) [0.070]	0.0721* (0.036) [0.070]	0.0390 (0.025) [0.100]	0.0610 (0.046) [0.220]	0.0564 (0.039) [0.160]	0.0564 (0.039) [0.160]	0.0479** (0.022) [0.030]	0.0495 (0.051) [0.360]	0.0282 (0.057) [0.610]	0.0282 (0.058) [0.610]	-0.0944 (0.124) [0.420]	0.0106 (0.067) [0.890]
Observations	1,376	1,376	1,376	1,333	1,376	1,376	1,376	1,333	1,376	1,376	1,376	1,333
R-squared	0.246	0.997	0.998	0.991	0.262	0.997	0.999	0.991	0.229	0.979	0.982	0.935
Indicator for 7-Eleven	X				X				X			
State fixed effects		X	X	X		X	X	X		X	X	X
State-specific trends			X				X				X	
Excluding DF				X				X				X
Mean dep. var.	0.82	0.82	0.82	0.41	1.32	1.32	1.32	0.68	4.61	4.61	4.61	2.16

Notes: This table shows the main results from expanding 7-Eleven access before, during, and after the media campaign. Observations are at the state-month level. Voluntary savers are the total number of accounts in a state-month with at least one voluntary contribution (measured in thousands). Voluntary contributions are the total number of contributions to voluntary accounts (measured in thousands). Amount contributed is the total amount from voluntary contributions (measured in millions of Mexican pesos). The outcome variables are transformed into logs for the estimation. The mean of the dependent variable for the treatment municipalities at baseline is shown. Regressions include month-year FE, and are weighted by the number of accounts pre-treatment. Robust standard errors clustered by state in parentheses. Wild cluster bootstrap p-values in brackets (stars denote significance from these values).

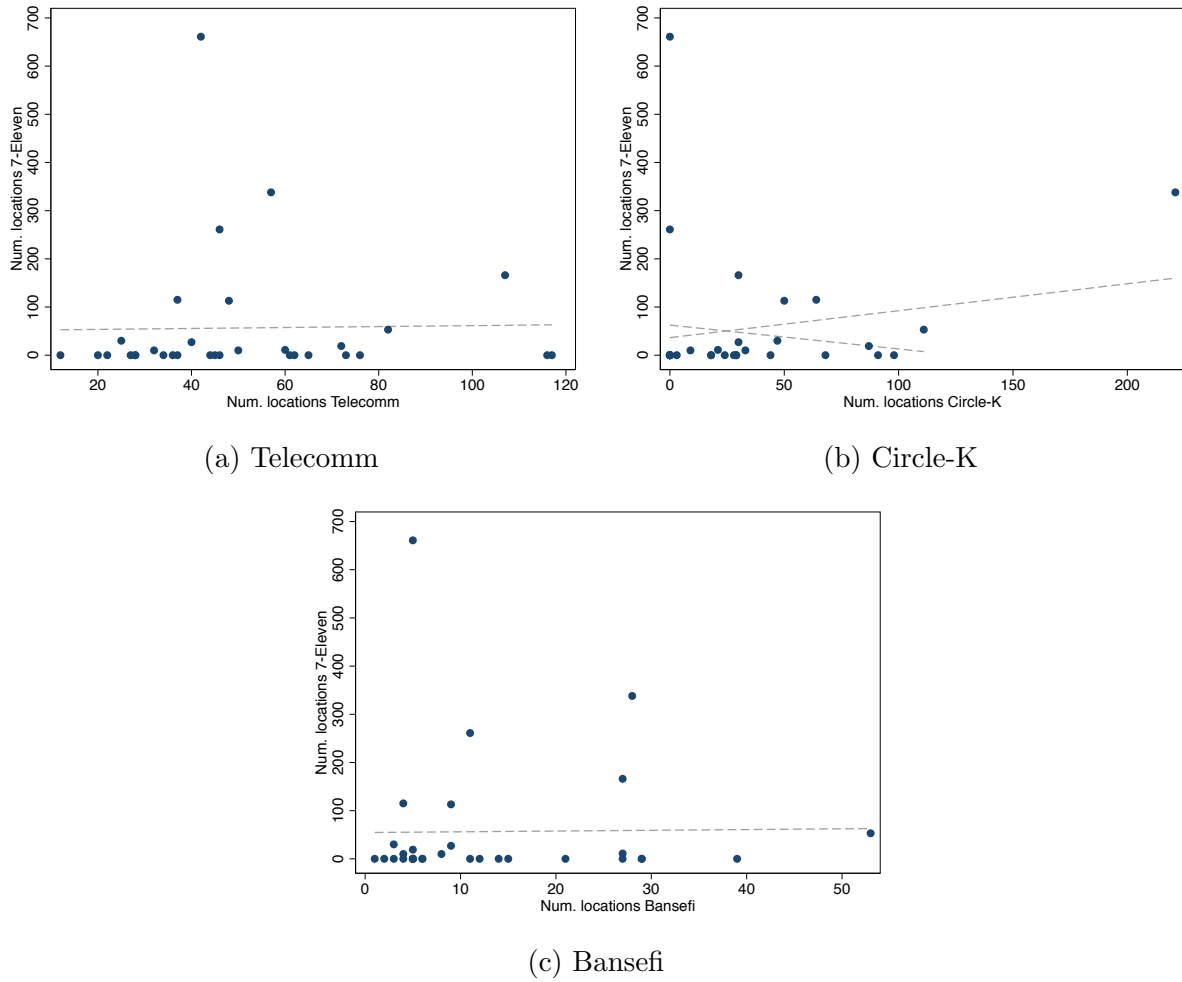
*** p<0.01, ** p<0.05, * p<0.1

Figure A1:
Still images from the ad campaign on television



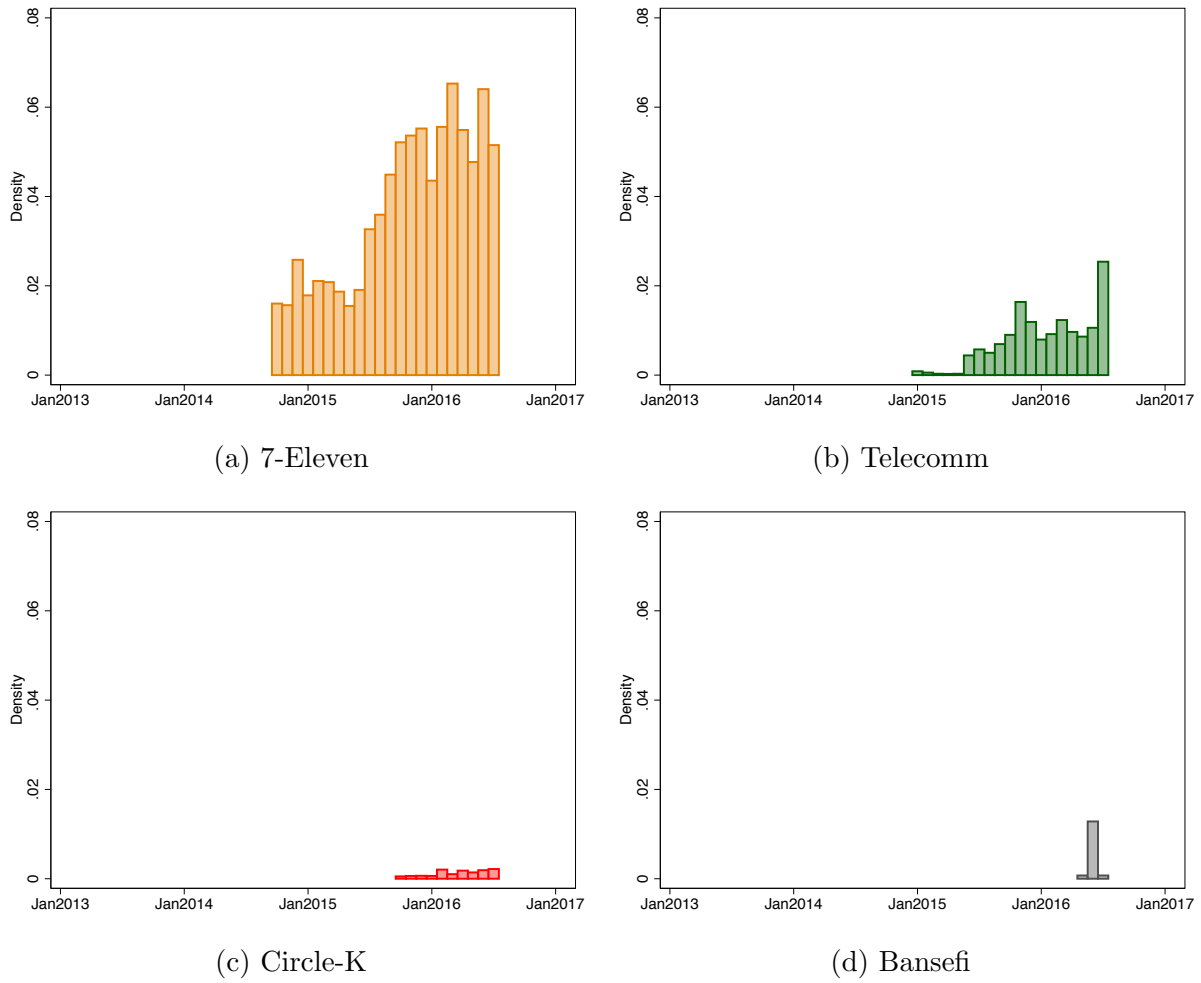
Notes: These four images are representative of the message the ad campaign was transmitting. The top left image shows an individual saving in 10 peso coins, with the objective of depositing the money in their AFORE. The top right image shows the individual going to a 7-Eleven convenience store. The image on the bottom left shows the worker making a voluntary contribution, and the bottom right image emphasizes the 7-Eleven policy. These stills were taken from a version of the ad with English subtitles, available at <https://www.youtube.com/watch?v=uSdOpwVJy1o> (last accessed April 4, 2018).

Figure A2:
Relationship between 7-Eleven state presence and other stores or
institutions




Notes: These plots show the relationship between the number of 7-Eleven stores by state and the number of locations for Telecomm, Circle-K, and Bansefi. The dashed lines represent the linear fit. In the second panel, the upward sloping linear fit considers all the data for 32 states, while the downward sloping line excludes the outlier (representing Distrito Federal).

Figure A3:
Histograms of voluntary contributions by store or institution




Notes: These plots show histograms detailing the density of voluntary contributions at each chain of stores or financial institution over time.

Figure A4:
7-Eleven presence by state over time




(a) Number of locations




(b) Change in locations

Notes: These plots show 7-Eleven locations by state pre and post-treatment (2014 vs 2016). The top graph shows the histograms, while the bottom graph plots the percentage change in number of locations over time, for each of the 32 states.

Figure A5:
Distribution of workers' state of birth and state of residence



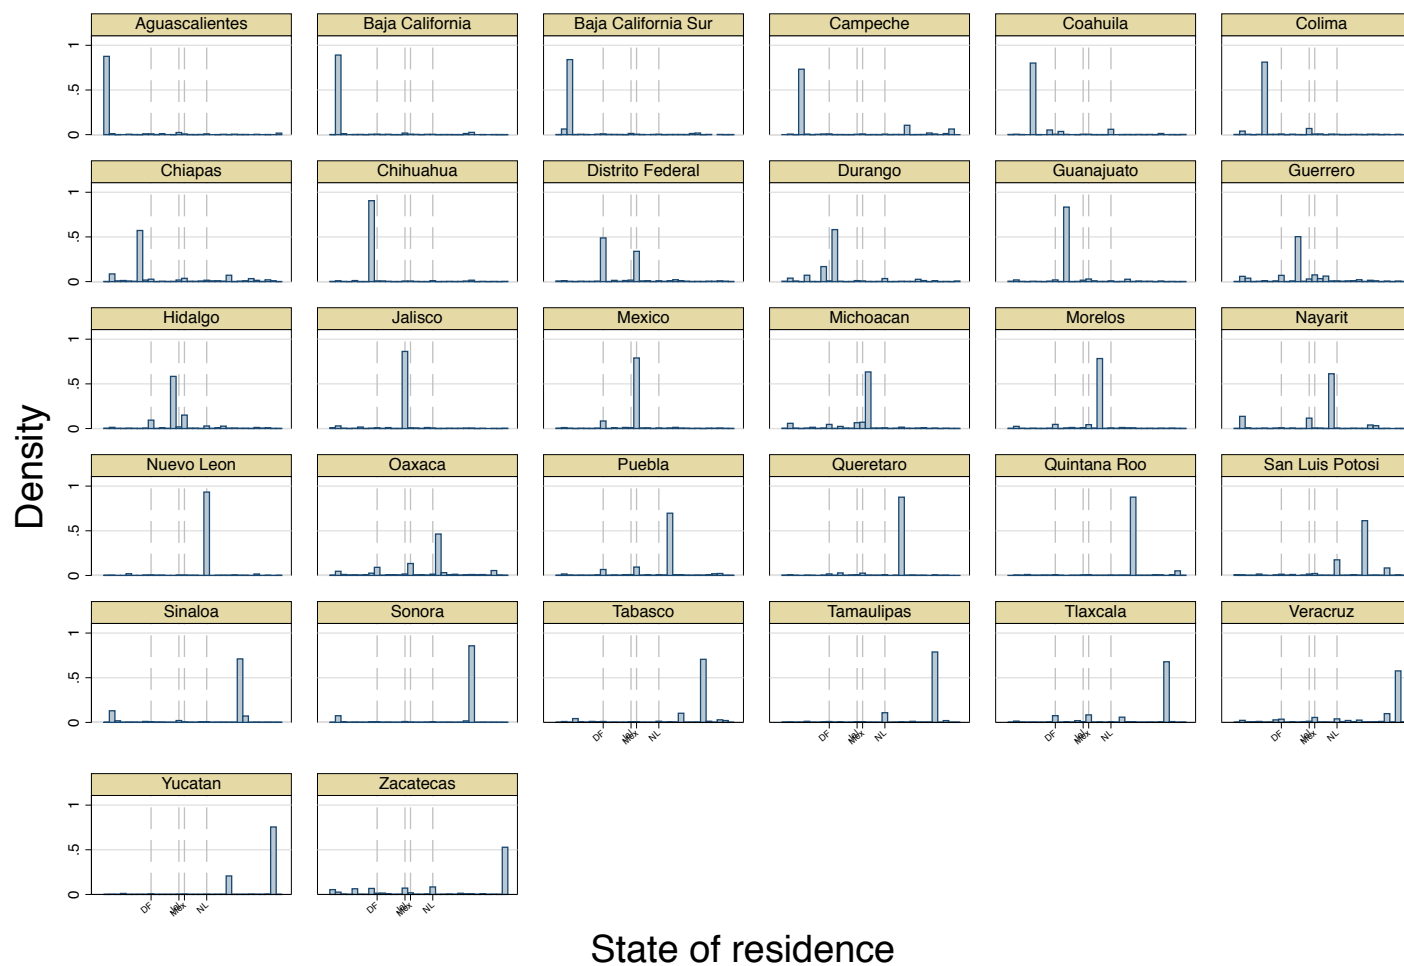
(a) Distribution of state of birth



(b) Distribution of state of residence

Notes: These plots show histograms for workers' state of birth and state of residence (measured in the last quarter of 2015).


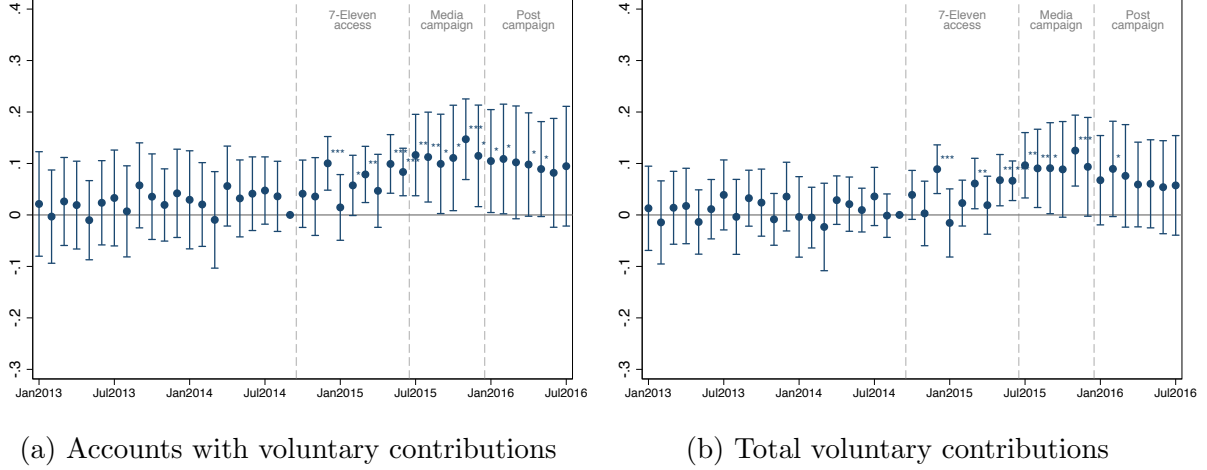
Figure A6:
Relationship between workers' state of birth and state of residence



Graphs by state of birth

Notes: These plots show the relationship between state of birth and state of residence (measured in the last quarter of 2015). Dashed lines correspond to states that contain the three largest metropolitan areas: Mexico City (Distrito Federal [DF] and Mexico [Mex]), Guadalajara (Jalisco [Jal]), and Monterrey (Nuevo Leon [NL]).


Figure A7:
Effect of treatments on voluntary contributions: State-level
estimates using state of birth



Notes: These plots show the main results from expanding 7-Eleven access before, during, and after the media campaign, within a dynamic DD framework. Observations are at the state-month level. The first panel shows the total number of accounts in a state-month with at least one voluntary contribution (measured in thousands), the second shows the total number of contributions to voluntary accounts (measured in thousands), and the third corresponds to the total amount from voluntary contributions (measured in millions of Mexican pesos). The outcome variables are transformed into logs for the estimation. Coefficients for month indicators interacted with 7-Eleven presence are shown, from regressions that include state and month-year FE, state-specific quadratic trends, and that are weighted by the number of accounts pre-treatment. Bars correspond to 95% confidence intervals based on robust standard errors clustered by state. Stars denote significance using wild cluster bootstrap p-values.


*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Figure A8:
Histograms of state measures of 7-Eleven intensity



(a) Number of stores

(b) Stores per 100,000 workers



(c) Market share

Notes: These plots show histograms displaying the frequency for different measures of 7-Eleven intensity of exposure for the 2,298 municipalities in the analysis. The first panel considers the total number of stores, the second panel shows the number of 7-Eleven stores per 100,000 active worker accounts (as measured prior to October 2014), and the third depicts 7-Eleven's market share (the number of 7-Eleven stores as a fraction of all convenience stores in the municipality). Note that these plots exclude the municipalities in which 7-Eleven has no market presence. The second panel also omits an outlier of 4,545 stores per 100,000 accounts in one particular municipality.

B Effects on Labor Market Outcomes

This section uses the National Occupation and Employment Survey (ENOE) to analyze effects on labor market outcomes. The ENOE is a nationally-representative survey with four rounds per year. All employment variables are directed at individuals ages 15 and older. We generate municipality-level aggregates for each survey round, detailing the share of individuals in a municipality that work and the share that work in the formal sector. We then condition on individuals that are currently working and calculate average monthly income at the municipality level, and the average number of hours worked per week.

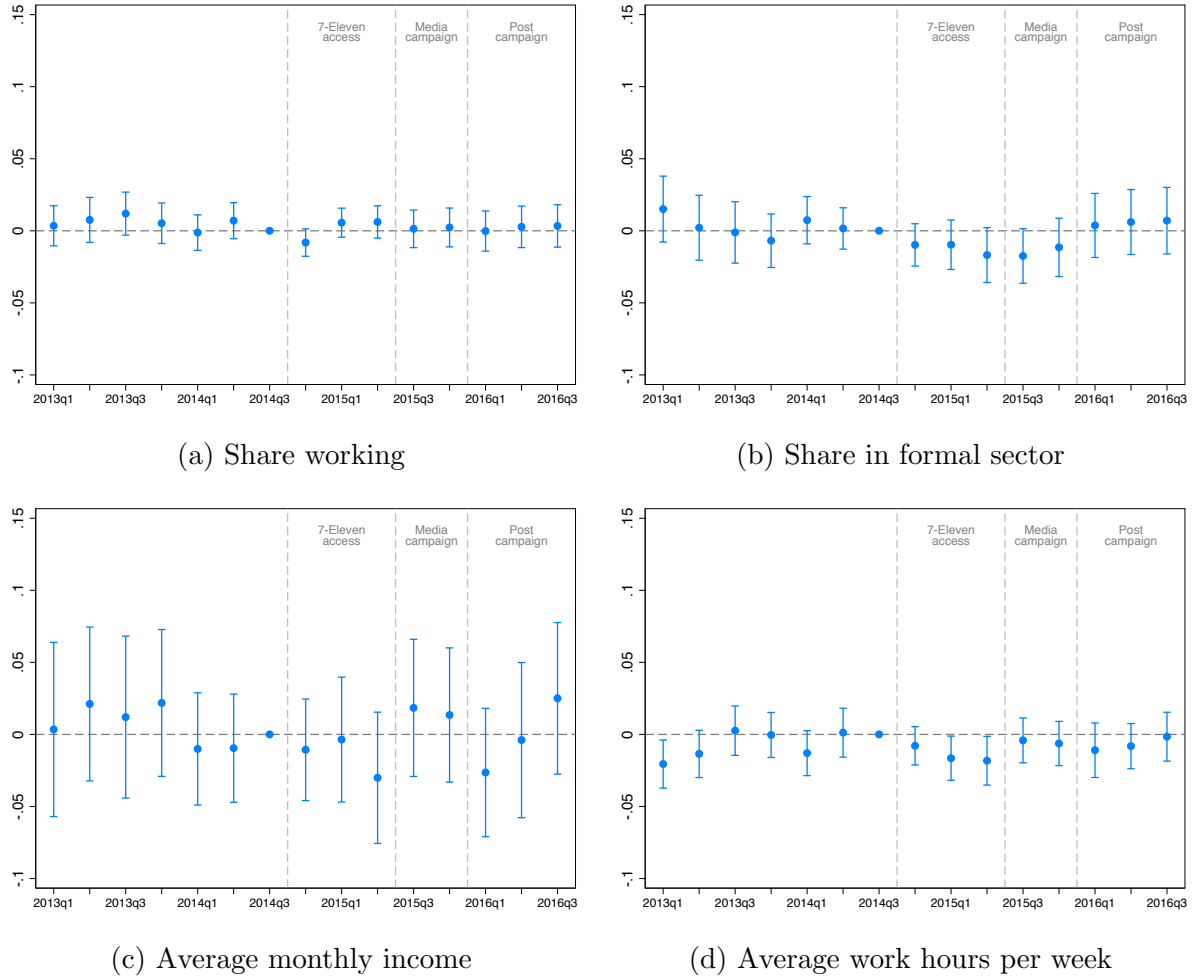
We estimate an equation similar to our dynamic DD of the form:

$$y_{mt} = \sum_{k=1}^T \beta_k (\mathbb{1}_{[7\text{-}Eleven]_m} \times \mathbb{1}_{[t=k]}) + \gamma_m + \theta_t + \nu_{mt} \quad (\text{B1})$$

where t denotes a quarter-year, y_{mt} is a municipality-quarter labor market outcome, and everything else is as defined in the main text.

Figure B1 shows the results from estimating equation B1 on our panel of municipality-quarters from the ENOE. We use the third quarter of 2014 as our excluded period. Each graph corresponds to one of the four outcomes. Error bars for the 95% confidence interval are shown, using standard errors clustered at the municipality level. Each graph shows estimates that are insignificant at the usual levels. This indicates that our treatments are not affecting any other labor market outcomes, which could confound our main results.

Figure B1:
Effect of treatments on voluntary contributions



Notes: These plots show effects on labor supply outcomes using the same time periods from expanding 7-Eleven access before, during, and after the media campaign, within a dynamic DD framework. Observations are at the municipality-quarter level, using data from the quarterly ENOE surveys. The outcomes are the share of individuals in a municipality-quarter that are working, the share that are working in the formal sector, the average monthly income conditional on working, and the average number of hours worked per week conditional on working. The last two outcomes are measured in logs. Coefficients for month indicators interacted with 7-Eleven presence are shown, from regressions that include municipality and quarter-year FE. Bars correspond to 95% confidence intervals based on robust standard errors clustered by municipality.

C Intensive and Extensive Margins

We partition the data to explore effects on the intensive and extensive margins. On the intensive margin, we consider individuals who made at least one voluntary contribution prior to October 2014. On the extensive margin, we consider those that did not have any voluntary savings prior to the treatment. We calculate the municipality aggregates and estimate equation 1.

Table C1 presents the results. Odd-numbered columns consider the intensive margin, and even-numbered columns the extensive one. The results for the intensive margin show significant estimates for voluntary savers and contributions, although the magnitudes are smaller than the main results in Table 3. There is evidence of a persistent effect after the media campaign ended, although at a smaller magnitude than the effect during the campaign. Results for amount contributed are all insignificant.

Columns 2, 4 and 6 indicate large effects on the extensive margin. Given that the outcomes were all zero prior to the treatment, the magnitudes of the estimates may be misleading, since they must be positive and large by construction. Note however that the magnitudes indicate that the strongest effect was experienced during the campaign, with similar magnitudes in the post-campaign period.

Overall, these results suggest that the policies had an effect both on individuals who were already making voluntary contributions and those that were not. This allows us to conclude that these policies are effective not only at inducing savings in non-savers, but also at increasing savings rates among savers.

Table C1:
Effect of treatments on voluntary contributions: Intensive and
extensive margins

	Voluntary savers		Voluntary contributions		Amount (in thousands of pesos)	
	(1)	(2)	(3)	(4)	(5)	(6)
7-Eleven \times access only	0.0433*** (0.016)	1.782*** (0.203)	0.0331* (0.017)	1.896*** (0.218)	-0.0263 (0.037)	2.530*** (0.299)
7-Eleven \times media campaign	0.0649*** (0.019)	2.012*** (0.221)	0.0744*** (0.020)	2.128*** (0.241)	-0.0001 (0.047)	2.603*** (0.282)
7-Eleven \times post-campaign	0.0439** (0.021)	2.085*** (0.226)	0.0492** (0.023)	2.184*** (0.244)	0.0133 (0.044)	2.542*** (0.273)
Observations	55,341	98,685	55,341	98,685	55,341	98,685
R-squared	0.991	0.893	0.988	0.896	0.937	0.863
Accts. with VCs pre-treatment	X		X		X	
Accts. without VCs pre-treatment		X		X		X
Coefficient tests:						
$H_0 : \beta_1 = \beta_2$	0.03	0.00	0.00	0.00	0.54	0.21
$H_0 : \beta_2 = \beta_3$	0.05	0.00	0.05	0.02	0.75	0.15
Mean dep. variable	113.93	0.00	184.52	0.00	656.81	0.00

Notes: This table shows the main results from expanding 7-Eleven access before, during, and after the media campaign, splitting the data by accounts with and without voluntary contributions (VCs) prior to the treatment (odd and even columns, respectively). Observations are at the municipality-month level. Voluntary savers are the total number of accounts in a municipality-month with at least one voluntary contribution. Voluntary contributions are the total number of contributions to voluntary accounts. Amount contributed is the total amount from voluntary contributions (measured in thousands of Mexican pesos). The outcome variables are transformed into logs for the estimation. The mean of the dependent variable for the treatment municipalities at baseline is shown. Regressions include municipality and month-year FE, and are weighted by the number of accounts pre-treatment. Robust standard errors clustered by municipality in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

D Heterogeneous Effects of the Main Results

We explore relevant heterogeneity in the main results based on the characteristics that we observe. We stratify the data by gender, age groups, year of affiliation, pre-treatment balance of the main retirement account, and by AFORES' initial voluntary savings rate. For each subsample, we calculate the relevant municipality-level aggregates, and estimate equation 1 as before.

We find that the treatment effects are stronger for men, younger workers, workers that affiliated with the system after 1997, workers with lower retirement account balances, and workers affiliated with fund managers with lower pre-treatment voluntary savings rates. These results are consistent with the idea that men are more likely than women to work formal jobs, that both younger and recently enrolled individuals depend on their individual accounts for retirement (note that older enrollees will retire under the previous defined benefits plan, even though they have their individual account under the new system), and that those with lower savings have more to gain.

Gender

The summary statistics in Table 1 of the main text show that around 36% of the 19 million active worker accounts are women, with a slightly larger proportion in the sample of workers that made at least one voluntary contribution between 2013 and 2016. This is consistent with women being less present in the labor force, but also with women being (slightly) more likely to work in informal sector jobs.¹

Table D1 presents the results separately by gender. Results for women are generally smaller in magnitude and less significant than those for men. In terms of the first two outcomes that we explore, for women we cannot reject that the magnitude of the effects during and after the media campaign are the same, while for men we obtain significantly

¹For example, according to the National Occupation and Employment Survey (ENOE) carried out by the National Institute of Statistics (INEGI), in the first trimester of 2013, 30.3% of women were employed in the informal market relative to 27.6% of men.

larger magnitudes post-campaign. The last two columns show that there are no statistically significant effects for women in terms of amount contributed voluntarily, while for men there is a significant 10% increase during the media campaign.

Age

Table D2 presents the results stratifying by age groups. We divide the sample into four categories based on a worker's age in January 2013: individuals younger than 30, workers between 30 and 49 years old, those between 50 and 64, and workers aged 65 and over. Our rationale is that the first group are very young workers, the second group are workers in their prime, the third group are those nearing retirement, while the last group are those that could or should have already retired.

We present the estimates from equation 1, using municipality aggregates based on these four groups. Columns 1 through 4 in Table D2 show that the main results on the total number of voluntary savers are driven by the youngest workers under 30 years old, although results are also significant and sizable for workers between 30 and 64 years old. The same holds for the total number of voluntary contributions in columns 5 to 8. The last four columns show that there are significant changes in the amount contributed voluntarily only for the youngest workers. Note that in general the access-only coefficients are not significant (except for the youngest group), and that there are strong persistent effects in the post-campaign months.

Year of Affiliation

Table D3 stratifies the sample based on workers' year of affiliation. Recall that year of affiliation refers to the first time that the worker registered with the social security system, oftentimes highly correlated with the first time the worker entered the formal labor force. Workers in the informal market need not affiliate. Since the new pension system based on individual retirement accounts began in 1997, we classify workers as those that registered

with the social security system before and after 1997. Figure D1a shows the distribution for the full population. All workers that registered in 1997 and after will retire under the new individual accounts system. However, workers that registered prior to 1997, with a few caveats, will mostly retire under the prior pay-as-you-go system, even though they possess an individual retirement account.

The results in Table D3 show that the effects are largest for workers that affiliated in 1997 or after. All coefficients corresponding to this group are positive, large and significant. There is also strong evidence of a persistent effect post-campaign of a larger magnitude than the effect experienced throughout the campaign months. The estimates for workers that affiliated prior to 1997 are about half the magnitude of the other workers, and are only significant during and after the campaign (and completely insignificant when focusing on the amount contributed).

This result is consistent with the incentives given the retirement plan that different workers face. Workers retiring under the pay-as-you-go system do not depend on their own savings account for their pension. However, note that the main results are not all driven by the more recent affiliates, meaning that even workers retiring under the old system found it beneficial to increase their savings. Although there is a correlation between age and year of affiliation, tying these results to those in Table D2, there is still considerable variation in a worker's age by year of affiliation. Figure D1b shows the relationship between year of affiliation and age.

Account Balance

Next, we present separate results by quartile of the pre-treatment balance of the main retirement account. This classification is based on the full distribution of the pre-treatment main balance, using all workers and adjusting with sampling weights. Figure D2 shows the full distribution. This measure captures a worker's number of years contributing to the system (that is, time in the formal economy), labor income, and base salary reported to the government.

Note that although a negative correlation may exist between year of affiliation and number of years contributing, this is not a purely mechanical relationship, since many workers drop out of the formal labor force. Consider also that even though a positive relationship may exist between labor income and contribution salary, employers may be reporting different shares of total wages for the contribution salary.

Table D4 shows the results from this exercise. For the first two outcomes (number of voluntary savers and voluntary contributions), the results are largest for the first and second quartiles. In both cases, all three estimates are significant and the magnitude of the effect is growing from one period to the next. The third and fourth balance quartiles also show significant effects for the campaign and post-campaign periods, with significantly larger point estimates in the latter period. Coefficients for the amount contributed are less significant, although a similar pattern holds with the largest effects occurring for accounts in the first balance quartile.

AFORE Voluntary Savings

For our last exploration of heterogeneity, we consider the initial voluntary savings rate of workers' AFORES. First, for each worker we assign the first AFORE that we observe in our dataset. Note that we only observe the AFORE for each date that workers make voluntary transactions. For the set of workers with voluntary transactions prior to the treatment, we do not observe any changes in AFORE during the pre-treatment period for 94% of them. For the remaining 6%, we assign the first AFORE we observe between January 2013 and September 2014. For the set of workers that only began having voluntary savings after the treatment, we assign the first AFORE we observe, even if this may not be the AFORE they began with in the pre-treatment period. Lastly, for workers without any voluntary contributions we directly observe their AFORE in the pre-treatment period for 99.5% of our sample.

We then use this assignment of AFORES (mostly pre-treatment) as well as the number of workers with at least one voluntary contribution before October 2014 to construct the

share of workers for each AFORE that had voluntary savings pre-treatment. Figure D3 in the online appendix shows the distribution of this variable across AFORES. We use the median of this measure to classify workers into those that were registered with an AFORE that had high or low initial voluntary savings rates. We then obtain municipality-AFORE aggregates and estimate a regression similar to equation 1, adding AFORE fixed effects.

Table D5 presents these results. All estimates are positive and significant. However, coefficients for workers registered with an AFORE with below median voluntary savings prior to treatment are three to four times larger than the estimates for workers in AFORES with high initial voluntary savings rates. This holds across all three relevant outcomes. Furthermore, we only find evidence of a persistent effect post-campaign (that is larger in magnitude than the campaign effect) for workers in AFORES with low initial voluntary savings.


Table D1:
Effect of treatments on voluntary contributions: Heterogeneous
effects by gender

	Voluntary savers		Voluntary contributions		Amount (in thousands of pesos)	
	(1)	(2)	(3)	(4)	(5)	(6)
7-Eleven \times access only	0.0374 (0.030)	0.0664*** (0.022)	0.0111 (0.031)	0.0528** (0.021)	0.0535 (0.061)	0.0349 (0.040)
7-Eleven \times media campaign	0.105*** (0.035)	0.155*** (0.027)	0.0829** (0.041)	0.144*** (0.027)	0.0654 (0.058)	0.103*** (0.039)
7-Eleven \times post-campaign	0.125*** (0.043)	0.202*** (0.034)	0.0860* (0.050)	0.179*** (0.035)	0.0937 (0.062)	0.0328 (0.048)
Observations	90,257	96,879	90,257	96,879	90,257	96,879
R-squared	0.982	0.984	0.979	0.981	0.910	0.926
Sample	Women	Men	Women	Men	Women	Men
Coefficient tests:						
$H_0 : \beta_1 = \beta_2$	0.00	0.00	0.00	0.00	0.85	0.11
$H_0 : \beta_2 = \beta_3$	0.33	0.00	0.90	0.07	0.64	0.08
Mean dep. variable	46.92	64.30	76.76	103.36	191.17	450.01


Notes: This table shows the main results from expanding 7-Eleven access before, during, and after the media campaign, stratifying by workers' gender. Observations are at the municipality-month level. Voluntary savers are the total number of accounts in a municipality-month with at least one voluntary contribution. Voluntary contributions are the total number of contributions to voluntary accounts. Amount contributed is the total amount from voluntary contributions (measured in thousands of Mexican pesos). The outcome variables are transformed into logs for the estimation. The mean of the dependent variable for the treatment municipalities at baseline is shown. Regressions include municipality and month-year FE, and are weighted by the number of accounts pre-treatment. Robust standard errors clustered by municipality in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Figure D1:
Workers' year of affiliation




(a) Distribution



(b) Relationship with age


Notes: The plot on the top shows the density of the year of affiliation of workers. All workers are included, using survey weights. The dashed line shows 1997, the year when the new individual retirement accounts system began. The plot on the bottom shows the share of workers affiliated in 1997 or after by age.

Figure D2:
Distribution of pre-treatment main account balance



Notes: This plot shows the density of the main account balance of workers prior to October 2014, measured in logs. All workers are included in the solid line, using survey weights. The dashed line considers only the workers that ever made a voluntary contribution, regardless of the timing of their first contribution.

Figure D3:
Voluntary contributions by AFORE pre-treatment



(a) Rate of early savers

(b) Distribution of early savers rate

Notes: The plot on the left shows the number of early savers (before October 2014) per 1,000 worker accounts in the pre-treatment period by AFORE. The plot on the right shows a histogram for these data.

Table D2:
Effect of treatments on voluntary contributions: Heterogeneous
effects by age

	Voluntary savers				Voluntary contributions				Amount (in thousands of pesos)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
7-Eleven × access only	0.151*** (0.041)	0.0383* (0.022)	0.0308 (0.024)	0.00719 (0.043)	0.115*** (0.042)	0.0211 (0.022)	0.0273 (0.026)	0.0133 (0.048)	0.281*** (0.074)	0.0316 (0.041)	0.0340 (0.059)	-0.0763 (0.123)
7-Eleven × media campaign	0.248*** (0.045)	0.121*** (0.027)	0.103*** (0.030)	-0.0300 (0.042)	0.221*** (0.050)	0.103*** (0.030)	0.111*** (0.034)	-0.0198 (0.051)	0.195*** (0.061)	0.0869** (0.042)	0.125*** (0.059)	-0.0719 (0.103)
7-Eleven × post-campaign	0.339*** (0.057)	0.154*** (0.034)	0.138*** (0.036)	-0.0871 (0.058)	0.285*** (0.064)	0.125*** (0.036)	0.138*** (0.040)	-0.0751 (0.068)	0.328*** (0.065)	0.0663 (0.048)	0.0436 (0.066)	-0.205 (0.145)
Observations	94,686	93,826	79,206	47,687	94,686	93,826	79,206	47,687	94,686	93,826	79,206	47,687
R-squared	0.964	0.986	0.978	0.906	0.959	0.983	0.975	0.902	0.878	0.927	0.894	0.763
Age group	< 30	30-49	50-64	> 64	< 30	30-49	50-64	> 64	< 30	30-49	50-64	> 64
Coefficient tests:												
$H_0 : \beta_1 = \beta_2$	0.00	0.00	0.00	0.23	0.00	0.00	0.00	0.38	0.23	0.23	0.12	0.97
$H_0 : \beta_2 = \beta_3$	0.00	0.05	0.08	0.18	0.05	0.24	0.24	0.24	0.02	0.64	0.13	0.27
Mean dep. variable	15.08	65.11	27.14	4.00	25.48	108.94	40.83	5.04	45.22	270.06	263.52	63.72

Notes: This table shows the main results from expanding 7-Eleven access before, during, and after the media campaign, stratifying by age groups. Observations are at the municipality-month level. Voluntary savers are the total number of accounts in a municipality-month with at least one voluntary contribution. Voluntary contributions are the total number of contributions to voluntary accounts. Amount contributed is the total amount from voluntary contributions (measured in thousands of Mexican pesos). The outcome variables are transformed into logs for the estimation. The mean of the dependent variable for the treatment municipalities at baseline is shown. Regressions include municipality and month-year FE, and are weighted by the number of accounts pre-treatment. Robust standard errors clustered by municipality in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table D3:
Effect of treatments on voluntary contributions: Heterogeneous
effects by year of affiliation

	Voluntary savers		Voluntary contributions		Amount (in thousands of pesos)	
	(1)	(2)	(3)	(4)	(5)	(6)
7-Eleven \times access only	0.0133 (0.023)	0.0917*** (0.027)	0.00132 (0.022)	0.0605** (0.027)	0.00536 (0.045)	0.115** (0.054)
7-Eleven \times media campaign	0.0780*** (0.028)	0.180*** (0.033)	0.0671** (0.030)	0.154*** (0.036)	0.0684 (0.044)	0.128*** (0.049)
7-Eleven \times post-campaign	0.103*** (0.034)	0.231*** (0.042)	0.0780** (0.036)	0.198*** (0.047)	0.0546 (0.049)	0.130** (0.053)
Observations	88,021	96,320	88,021	96,320	88,021	96,320
R-squared	0.986	0.980	0.983	0.977	0.925	0.915
Affiliation year	< 1997	\geq 1997	< 1997	\geq 1997	< 1997	\geq 1997
Coefficient tests:						
$H_0 : \beta_1 = \beta_2$	0.00	0.00	0.00	0.00	0.19	0.79
$H_0 : \beta_2 = \beta_3$	0.13	0.01	0.56	0.06	0.73	0.97
Mean dep. variable	71.89	39.33	115.24	64.88	499.36	141.81

Notes: This table shows the main results from expanding 7-Eleven access before, during, and after the media campaign, stratifying by year of affiliation. Observations are at the municipality-month level. Voluntary savers are the total number of accounts in a municipality-month with at least one voluntary contribution. Voluntary contributions are the total number of contributions to voluntary accounts. Amount contributed is the total amount from voluntary contributions (measured in thousands of Mexican pesos). The outcome variables are transformed into logs for the estimation. The mean of the dependent variable for the treatment municipalities at baseline is shown. Regressions include municipality and month-year FE, and are weighted by the number of accounts pre-treatment. Robust standard errors clustered by municipality in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table D4:
Effect of treatments on voluntary contributions: Heterogeneous
effects by pre-treatment balance of main account

	Voluntary savers				Voluntary contributions				Amount (in thousands of pesos)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
7-Eleven × access only	0.177*** (0.030)	0.132*** (0.035)	0.0492 (0.038)	0.0371 (0.025)	0.212*** (0.038)	0.0989** (0.040)	0.0304 (0.038)	0.0221 (0.024)	0.108 (0.074)	0.0667 (0.079)	-0.00706 (0.062)	0.0145 (0.046)
7-Eleven × media campaign	0.292*** (0.032)	0.222*** (0.037)	0.135*** (0.044)	0.107*** (0.031)	0.371*** (0.042)	0.185*** (0.046)	0.120** (0.049)	0.0913*** (0.033)	0.174** (0.075)	0.0805 (0.090)	0.0713 (0.069)	0.0771* (0.046)
7-Eleven × post-campaign	0.397*** (0.044)	0.335*** (0.045)	0.214*** (0.054)	0.132*** (0.037)	0.464*** (0.058)	0.290*** (0.054)	0.191*** (0.061)	0.113*** (0.039)	0.166* (0.089)	0.132* (0.076)	0.112 (0.081)	0.0641 (0.048)
Observations	93,955	88,494	84,925	80,109	93,955	88,494	84,925	80,109	93,955	88,494	84,925	80,109
R-squared	0.945	0.956	0.969	0.988	0.938	0.949	0.965	0.986	0.839	0.832	0.865	0.924
Balance quartile	1	2	3	4	1	2	3	4	1	2	3	4
Coefficient tests:												
$H_0 : \beta_1 = \beta_2$	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.00	0.40	0.87	0.19	0.21
$H_0 : \beta_2 = \beta_3$	0.00	0.00	0.00	0.07	0.02	0.00	0.01	0.21	0.92	0.51	0.57	0.72
Mean dep. variable	11.09	11.22	18.80	70.11	14.44	17.21	30.92	117.55	98.90	49.69	79.41	413.17

Notes: This table shows the main results from expanding 7-Eleven access before, during, and after the media campaign, stratifying by quartiles of the main account balance prior to treatment. Observations are at the municipality-month level. Voluntary savers are the total number of accounts in a municipality-month with at least one voluntary contribution. Voluntary contributions are the total number of contributions to voluntary accounts. Amount contributed is the total amount from voluntary contributions (measured in thousands of Mexican pesos). The outcome variables are transformed into logs for the estimation. The mean of the dependent variable for the treatment municipalities at baseline is shown. Regressions include municipality and month-year FE, and are weighted by the number of accounts pre-treatment. Robust standard errors clustered by municipality in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table D5:
Effect of treatments on voluntary contributions: Heterogeneous
effects by initial voluntary savings rate of AFORE

	Voluntary savers		Voluntary contributions		Amount (in thousands of pesos)	
	(1)	(2)	(3)	(4)	(5)	(6)
7-Eleven \times access only	0.554*** (0.079)	0.182*** (0.028)	0.651*** (0.091)	0.192*** (0.032)	0.750*** (0.108)	0.309*** (0.057)
7-Eleven \times media campaign	0.768*** (0.099)	0.224*** (0.035)	0.888*** (0.112)	0.243*** (0.043)	0.854*** (0.109)	0.319*** (0.055)
7-Eleven \times post-campaign	0.989*** (0.110)	0.227*** (0.039)	1.124*** (0.122)	0.244*** (0.045)	1.089*** (0.109)	0.294*** (0.046)
Observations	494,070	592,884	494,070	592,884	494,070	592,884
R-squared	0.655	0.751	0.657	0.744	0.546	0.663
Below median AFORE VCs	X		X		X	
Above median AFORE VCs		X		X		X
Coefficient tests:						
$H_0 : \beta_1 = \beta_2$	0.00	0.01	0.00	0.02	0.02	0.78
$H_0 : \beta_2 = \beta_3$	0.00	0.87	0.00	0.98	0.00	0.51
Mean dep. variable	1.67	13.45	2.40	22.15	9.80	71.77

Notes: This table shows the main results from expanding 7-Eleven access before, during, and after the media campaign, stratifying by voluntary savings rates pre-treatment for each AFORE (above and below the median), considering only the individuals' pre-treatment AFORE. Observations are at the municipality-month level. Voluntary savers are the total number of accounts in a municipality-month with at least one voluntary contribution. Voluntary contributions are the total number of contributions to voluntary accounts. Amount contributed is the total amount from voluntary contributions (measured in thousands of Mexican pesos). The outcome variables are transformed into logs for the estimation. The mean of the dependent variable for the treatment municipalities at baseline is shown. Regressions include municipality, month-year, and AFORE (as registered in 2013) FE, and are weighted by the number of accounts pre-treatment. Robust standard errors clustered by municipality in parentheses.

*** p<0.01, ** p<0.05, * p<0.1