The Long-Term Effects of Temporary Incentives to Save: Evidence from a Field Experiment in Mexico

Paul Gertler
UC Berkeley
gertler@berkeley.edu

Sean Higgins
UC Berkeley
seanhiggins@berkeley.edu

Aisling Scott
UC Berkeley
aisling_scott@haas.berkeley.edu

Enrique Seira
ITAM
enrique.seira@itam.mx

August 20, 2017

Abstract

Are savings accounts experience goods, where consumers learn the value of saving in formal financial institutions only after opening and using an account? Across 110 bank branches throughout Mexico, we randomized a short-run incentive to save: prize-linked savings accounts with cash-prize lotteries, where lottery tickets are awarded as a function of savings balances. Both existing account holders and new account openers were eligible to win prizes, which were based on new savings accumulated over two months. After two months, the incentive was removed. We find that the savings lotteries served as a nudge to open accounts, causing a 43% increase in the number of bank accounts opened in treatment branches during the lottery months; in pre- and post-lottery months, there was no difference in the number of accounts opened in treatment and control branches. On the other hand, there was no intensive margin effect of the incentive on savings balances for existing accounts. While "compliers" who open accounts in treatment branches during lottery months initially save less than those who open accounts during the same months in control branches, we find that their savings balances catch up to the control group over time (after about 18 months). Furthermore, account openers in the treatment and control branches make transactions and keep their accounts active at similar rates in the five years after account opening.

1 Introduction

In developing countries, 55% of adults do not have a bank or mobile money account (Demirgüç-Kunt et al., 2015). Meanwhile, various studies have found positive impacts of access to savings, such as increased agricultural investment, harvest, and future consumption (Brune et al., 2016), increased microenterprise investment, profits, and future consumption (Dupas and Robinson, 2013), increased investment in childrens' education (Prina, 2015), reduced debt (Kast and Pomeranz, 2014), and reduced income volatility (Chamon et al., 2013). The poor face a number of barriers to opening accounts in financial institutions, however, such as a lack of access to banks (Burgess and Pande, 2005), transaction costs (Karlan et al., 2014), and low trust in banks (Bachas et al., 2017). A three-country experiment found that even after opening bank accounts, account use is low (Dupas et al., forthcoming).

To explore why the use of formal financial institutions is low despite the benefits found in the literature, we provide a temporary incentive to open and save in formal bank accounts by randomly offering the chance to win cash prizes in a lottery, where the number of lottery tickets an individual receives are a function of new savings generated over a two-month period. Specifically, among 110 branches of a government bank in Mexico (Bansefi), we randomly assigned 40 branches to be in a treatment group that advertized the lottery. The advertisements were done in the branch through posters as well as in the neighborhood surrounding the branch using flyers and loud-speaker cars, so that non-clients were also aware of the offer. Both existing clients and new clients who opened accounts were eligible for the prizes. For each 50 pesos of *new* savings generated during October and November 2010, a client received one lottery ticket. We then raffled one thousand small prizes of 400 pesos (US$32) and two large prizes of 10,000 pesos (US$809). After November 2010, the lotteries cease, so that the benefits of opening an account or saving in treatment and control branches no longer differ.

We find evidence consistent with the hypothesis that savings accounts are experience goods: the value of saving is not known until after the costs (both non-monetary and monetary if applicable) of opening an account and beginning to save are incurred. Under this hypothesis, the expected marginal benefit of opening an account may not be equal to the actual marginal benefit, and individuals who perceive the marginal benefit to be lower than the cost of opening an account will

not open an account. In any model, a temporary incentive that increases the marginal benefit of opening an account and beginning to save should lead to increased account openings—and we indeed observe a 43% increase in the number of accounts opened in treatment branches during lottery months, relative to control branches in the same months. In a standard model, however, once the temporary incentive is removed individuals for whom the marginal cost exceeds marginal benefit would no longer save in the account. We observe the opposite: although initially lower, the savings balances of "compliers" who open accounts in response to the incentive catch up to the savings balances of those who open accounts in control branches. The levels of active use of the account are similar across lottery month account openers in treatment and control branches.

Existing account holders (also eligible for lottery tickets based on each 50 pesos of *new* savings added to their account balances during the lottery months) would only benefit from saving more based on the increase in their expected income from the small chance of winning a lottery prize. Because they already have experience saving, there would not be any additional benefit beyond this increase in expected income; the increase in their expected utility from this can be thought of as analogous to a certainty-equivalent increase in the interest rate. It is thus consistent with the experience good hypothesis that we see no increase in saving on the intensive margin, especially given the low interest rate elasticities for saving found by Karlan and Zinman (2016).

Our experimental design differs from existing literature in several ways. First, we conduct the first randomized control trial (RCT) of prize-linked savings (PLS) accounts, which have previously been studied only in laboratory and lab-in-the-field experiments (Atalay et al., 2014; Filiz-Ozbay et al., 2015; Dizon and Lybbert, 2017). Second, rather than randomizing the offer of savings accounts at the individual level, we randomize at the branch level and allow individuals to self-select into opening accounts. Thus, we are studying a population likely at the margin of opening accounts.[1] Third, by testing for impact on both the extensive margin (account opening by new clients) and intensive margin (increases in saving by existing clients), as well as observing behavior of induced account openers after the lottery incentive ends, we shed light on the hypothesis of saving as an experience good. Fourth, we follow users for five years, making this one of the longest-run studies

[1] This approach also has a drawback, of course: the only experimental comparisons are those of the number of new accounts opened in treatment vs. control branches and the change in savings among existing account holders. We use comparisons of those who open accounts in treatment vs. control branches during lottery months—who are not comparable in the experimental sense—to study how those who are induced to open accounts by the lottery differ from standard account openers.

of savings in a developing country (with a notable exception being Suri and Jack, 2016).

Altough prize-linked savings accounts are a common financial product in places like the Latin America, the Middle East, Europe, the United Kingdom, and South Africa (Cole et al., 2007; Kearney et al., 2011), no study has evaluated their impact on savings using an RCT. Tufano et al. (2011) present survey evidence that PLS accounts might be particularly attractive for those who do not save; in a lab experiment, Atalay et al. (2014) find that offering PLS generates "new savers." Filiz-Ozbay et al. (2015) show theoretically and in a lab experiment that people who overweight small probabilities save more when offered a PLS account. Dizon and Lybbert (2017) replicate this result in a lab-in-the-field experiment in Haiti, and both Cole et al. (2016) and Dizon and Lybbert (2017) find that savings in PLS accounts do not substitute other savings, but that they crowd out gambling expenditures. This is consistent with Herskowitz's (2016) finding that gambling and savings technologies are both used to save for durables, and that there is substitution from gambling to saving when a secure savings device is provided. Policy interest in the impact of PLS accounts on savings—both among new and existing savers—has increased recently, in particular after they became legal in the US following the passage of the American Savings Promotion Act.

2 Context and Experimental Design

2.1 Banking in Mexico

Mexico's financial market is dominated by five large banks which had 73% market share in 2010 (Jiménez Bautista, 2012), and these banks have little interest in serving the poor. While Compartamos Banco has rapidly expanded access to microcredit (at high interest rates), they have not aggressively pursued microsavings (Angelucci et al., 2015).

In 2001, the Mexican government founded the National Savings and Financial Services Bank (Bansefi), with the mission of "contributing to the economic development of the country through financial inclusion... to strengthen savings and loans mainly for low income segments." Bansefi focused on fostering savings for the poor through low-cost savings accounts with no minimum balance. At the time of our experiment in 2010, Bansefi had 494 branches and about 5 million accounts.

2.2 Experimental Design

Sampling frame. Of the 494 Bansefi branches, the majority of them offered a match savings program with commitment device features, called *premiahorro*. We excluded these branches from the experiment since this type of account might be more attractive to new clients than an account with temporary lottery incentives. This left us with 214 branches for our sampling frame. We further restricted the sample by excluding the largest and smallest branches from our sample, measured by the volume of new accounts opened in the first half of 2010. We removed approximately the smallest 25% and largest 25% of branches from the sampling frame. In addition, we removed branches in any states that only contained one branch. After applying these selection criteria, our sampling frame consisted of 110 Bansefi branches spanning 19 of Mexico's 32 states throughout the entire country from Baja to the Yucatan Peninsula. One contribution of our paper is that experiments on savings rarely have this extent of geographical breadth.

Randomization. Within the 110 Bansefi branches in our sampling frame, we conducted a simple randomization to assign 40 branches to treatment. Appendix Table A1 shows that treatment and control branches have balanced covariates, both for overall characteristics about the branches as well as characteristics of the existing accounts at those branches. Figure 2 shows the locations of treatment and control branches.

Lotteries. In treatment branches, lotteries were announced beginning in mid-September 2010 through in-branch posters, as well as flyers and loud-speaker cars on nearby streets. Two lotteries were held, the first on October 12, 2010 and the second on November 12, 2010. For administrative reasons, only one type of account—*debicuenta* accounts or those tied to a debit card—were eligible; the other types of accounts offered by Bansefi (including a *cuentahorro* savings account not tied to a card) were not eligible. Both existing clients with a *debicuenta* account as well as new clients were eligible. Figure 1 shows the timeline of the experiment and an example of Bansefi's advertisements of the savings lotteries, which reads "save in a *debicuenta* account and multiply your money."

In each month preceding a lottery, debicuenta accountholders in treatment branches received one lottery ticket for each 50 pesos of *new* savings generated over the month; pre-existing balances were subtracted from the end balance to determine the amount of new savings generated. Other

than the lottery tickets, the other aspects of the account (including the interest rate) were identical to those of accounts in control branches. Furthermore, clients could deposit and withdraw their money at any time before, during, and after the lottery. Thus the lottery could be "gamed" by individuals depositing a large amount prior to the lottery and withdrawing it afterwards. We test for this type of behavior; although we find some limited evidence of gaming, it does not drive our results.

The lotteries were conducted at the national level (i.e., they were not stratified by branch). Each month, 1000 small prizes of 400 pesos (US$32) and two large prizes of 10,000 pesos (US$809) were awarded. The probability of winning is endogenous to total participation and was therefore not known ex ante. Ex post, in the October lottery the probability of a particular ticket winning was 1 in 713, and the median saver earned 27 tickets.

A substantial proportion of Bansefi's clients are beneficiaries of Mexico's large cash transfer program Oportunidades, who receive their benefits directly in Bansefi *debicuenta* accounts (Bachas et al., 2017). Oportunidades beneficiaries were also eligible for lottery prizes, but because their accounts are opened for them automatically by the government when they are enrolled in the program, we exclude Oportunidades accounts from the analysis.

3 Data

We use two types of administrative data from Bansefi for the accounts at the 110 branches included in our experiment. First, we have data on the number of accounts opeend each month in each branch from 2008 through May 2011. Second, we have transactions data for pre-existing accounts and those opened during lottery months over a five year period (longer in the case of pre-existing accounts). Specifically, for accounts opened during the lottery months, we have transactions data from the date they were opened through July 2015. For *debicuenta* accounts that already existed in treatment and control branches we have transactions data from 2008–2015.

4 Extensive Margin Effect of Savings Lotteries

We compare the average number of accounts opened in treatment and control branches. Separately for each month t we estimate the ANCOVA specification

$$y_j = \gamma T_j + \theta y_{j0} + \epsilon_j, \tag{1}$$

where y_j is the number of new non-Oportunidades debicuenta accounts opened at branch j in that month, T_j indicates that branch j is a treatment branch, and y_{j0} is the number of accounts opened per month at baseline. It is worth noting that the number of non-Oportunidades *debicuenta* accounts opened in a particular branch in a particular month is low; as discussed above, Bansefi also offers other types of accounts and a large proportion of Bansefi clients are Oportunidades beneficiaries whose accounts were created for them by the government cash transfer program. On average over our period, the number of new accounts opened in a particular branch in a particular month is only 3.6.

We estimate (1) for six months prior to the lotteries, the two months of the lotteries, and six months after the lotteries. We define the number of accounts opened per month at baseline as the average from all periods prior to the first month for which we estimate effects, i.e. prior to March 2010. The results (i.e. the γ coefficients for each month) are shown in Figure 3.

In non-lottery months (March 2010–September 2010 and December 2010–May 2011), we see no difference in the number of accounts opened in treatment and control localities. In the lottery months, October and November 2010, we see a statistically significant increase in the number of accounts opened. In October 2010, 1.6 additional accounts are opened, or 45% more than the number opened in control branches in the same month. Note that the control mean in that month is 3.6, statistically insignificant from the mean number of accounts opened in the control group over all periods, suggesting that the result is not driven by a substition effect where new account openers who would have opened accounts in control branches instead open them in treatment branches.

In November 2010, the effect is smaller, at 1.0 additional account or 40% relative to the mean in the control group that month. The lower effect is unsurprising, given that the second lottery occurred on November 12 and after this date, there was no incentive to open accounts. Averaging

effect sizes across the two lottery months, the lotteries induced a 43% increase in the number of accounts opened.

5 Intensive Margin Effects of Savings Lotteries

We use existing accounts at treatment and control branches to measure the effect of the lottery on the intensive margin. While the prizes changed the incentives to save which could affect savings levels among pre-existing account holders, any "experience good" effect where users learn the value of the accounts after being induced to open them by the lotteries would not be present for existing clients.

We measure savings using end-of-month balance, and again use an ANCOVA specification that controls for baseline savings balance. We again define baseline by averaging over early months; because we estimate effects up to one year before the lotteries (since October 2009), we restrict the analysis to accounts that had been opened prior to October 2009 so that baseline savings is defined. Finally, because clients in treatment branches may have won prizes from the lottery—which could in turn affect the amount they save in the account—we control for the amount won. Specifically, separately for each month t we estimate

$$Balance_{ij} = \gamma T_j + \psi Winnings_{ij} + \beta Balance_{ij0} + \epsilon_{ij} \tag{2}$$

where $Balance_{ij}$ is end-of-month balance of account i in branch j (in month t), T_j is a dummy that equals 1 for treatment branches, and we control for the amount won $Winnings_{ij}$ (which equals 0 in pre-lottery months, the amount won in October for the October 2010 regression, and the sum of winnings in the October and November lotteries for all months after October 2010), and $Balance_{ij0}$ is average baseline end-of-month balance.[2] Standard errors are clustered at the branch level.

The results—i.e. γ from (2) for each month t—are shown in Figure 4. There is no evidence of an intensive margin effect, as the coefficients for the months immediately following the lotteries are not statistically different from 0. For some later months (after about September 2011) there are some statistically significant but small coefficients (less than 200 pesos or US$16). Because these

[2]Following other papers measuring savings (e.g., de Mel et al., 2013; Dupas et al., 2016; Kast et al., 2016; Karlan and Zinman, 2016), we winsorize savings balances to avoid results driven by outliers. Our main results winsorize at the 95th percentile, and the results are robust to other cut-offs.

positive coefficients occur significantly after the lotteries and in some cases are not robust to other winsorizing cutoffs, we do not interpret this as an effect of the lotteries on savings at the intensive margin.

Other research has found that savings of the poor do not respond to changes in the interest rate (Karlan and Zinman, 2016), except at unsustainable interest rates on the order of 20% per year (Schaner, 2016). These findings are consistent with the lack of an intensive margin effect of the lotteries on savings among existing clients. On the other hand, the expected value interest on saving during lottery months is quite high. Although the clients would not have been able to calculate their expected value interest rate of saving during lottery months since the total number of tickets awarded was unknown ex ante, we can calculate the expected value ex post. Given that the chance of any ticket winning was 1 in 713 and the expected value of the prize conditional on winning was $(1000 \times 400 + 2 \times 10,000)/402 = 1045$ pesos, the expected winnings per 50 pesos saved was 1.5 pesos, or a 3% return over the month. Nevertheless, since the lottery incentive was temporary and based on *new* savings (i.e., additional savings from October did not increase expected "interest" in November), this 3% return over one month is not the same as a 3% monthly return.

6 Savings Accounts as an Experience Good

We now examine the behavior of those who are induced to open accounts by the lotteries. We compare those who open accounts in treatment branches in October 2010 to those who open accounts in the same month in control branches.[3] Two caveats are in order: first, this section is descriptive and seeks to explore how the behavior of those induced to open accounts by the lottery differ from other clients. Clearly, this comparison suffers from selection bias and does not tell us anything causal. Second, because the lotteries increased account openings in treatment branches by 45% in October 2010, about two-thirds of treatment branch lottery-monmth account openers are "always takers" who would have opened accounts anyway. Thus, if we are seeking to learn how "compliers" induced to open accounts by the lottery differ from others, our comparison will be biased toward zero.

[3]We focus on October 2010 since opening an account and saving was incentivized over the entire month; in November 2010 it was no longer incentivized after the November 12 lottery took place.

6.1 Savings Balances

We first compare average end-of-month balances each month for accounts opened in October 2010 in treatment branches vs. accounts opened in October 2010 in control branches. We estimate the following separately for each month t:

$$Balance_{ij} = \gamma T_j + \psi Winnings_{ij} + \epsilon_{ij} \tag{3}$$

where $Balance_{ij}$ is end of month balance of account i in branch j (in month t), T_j is a dummy that equals 1 for treatment branches, and we control for the amount won $Winnings_{ij}$ (summing over the October and November lotteries, except for t = October 2010 where $Winnings_{ij}$ is just the amount won in October). There is no control for baseline balance since these accounts did not exist prior to the lottery months. Standard errors are clustered at the branch level.

The γ coefficients for each month t are shown in Figure 5. Initially, we see that despite the attenuation bias described above, October account openers in treatment branches save substantially less than those who opened accounts in control branches in the same month. Over time, however, their savings catch up: after 18 months, the difference in savings levels of October openers in treatment and control branches are no longer statistically significant.

This finding is not due to a composition effect of "compliers" closing their accounts while "always takers" in treatment branches—who would have opened accounts in October 2010 anyway and are thus not distinguishable from lottery-month openers in control branches—leave their accounts open. While it is not uncommon to leave accounts dormant, these dormant accounts still enter into (3); only accounts that were formally closed could cause this. However, in April 2012, the first month in which savings of treatment branch lottery-month openers catch up to those from control branches, 95% of accounts opened in October 2010 are still open, and we cannot reject that there is no difference between treatment and control branches in any period. In June 2014, nearly four years after the lotteries and when we observe a tight zero in the difference in savings between treatment and control, 89% of accounts are still open. In no month can we reject that the proportion of accounts remaining open is equal across accounts opened in October 2010 in treatment and control branches.

6.2 Account Survival

Given the high rates of accounts remaining open described above, here we use a slightly stricter measure of account survival: the account must have at least 50 pesos (US$4). Figure 6 shows the proportion of accounts with at least a 50 peso end-of-month balance in treatment (blue squares) and control (orange circles). The proportion of accounts that are still used falls over time, but the levels and trend at which it falls are almost identical between accounts opened in October 2010 in treatment and control accounts.[4] One year after the lotteries, 75% of accounts remain in use by this measure; three years after, 47%; and nearly five years after, 35%. In all of these cases, we cannot reject the null hypothesis of no difference between lottery-month openers in treatment and control branches. Five years after receiving the card, this figure is around 20%.

6.3 Transactions

For accounts that remain in use, we examine active use of the account using three measures from the literature: at least one deposit in the last six months (Schaner, forthcoming), at least two deposits in the last six months (Dupas and Robinson, 2013), and the more long-term measure of at least five deposits in the last two years (Dupas et al., forthcoming). Figure 7 shows the proportion of accounts with a balance of at least 50 pesos that are actively used by these criteria.

Using the least restrictive measure of at least one deposit, about 42% of accounts are active users six months after opening the account in April 2011—the first month for which we measure this since opening an account requires making a deposit (and hence the proportion would be 100% in prior months from the October 2010 deposit). In most periods, there is no statistically significant difference in this measure; in the periods where there is a statistically significant difference, the treatment accounts use the account more. Using the slightly more restrictive measure of at least two deposits, about 22% are active users initially, and 12% are active users after five years. Finally, using Dupas et al.'s (forthcoming) longer-horizon measure, the proportion of active users is fairly constant around 20%, and again we cannot reject equal activity between treatment and control accounts.

[4]We graph conditional means to account for the possibility of winning the lottery: specifically, the conditional mean is the coefficient α from a regression of $\mathbb{I}(Balance_{ij} > 50) = \alpha + \psi Winnings_{ij} + \epsilon_{ij}$ separately for each month t, where $\mathbb{I}$ is the indicator function that equals 1 if its argument is true and 0 otherwise.

Figure 8 shows the distribution of deposits and withdrawals in surviving accounts in accounts opened during the lottery month in treatment branches (outlined in blue) and control branches (solid orange). While use falls over time, the distribution of use between the treatment and control group looks similar, and in both a group of very active users making 20 or more transactions per year (about 10% of clients) persists in the long term.

7 Conclusion

We find that prize-linked savings accounts can increase saving on the extensive margin by inducing new savers to open accounts. We thus validate in a randomized control trial conducted across Mexico a result suggested by Tufano et al. (2011) based on survey data and confirmed in the lab by Atalay et al. (2014). The lottery prizes were only offered over a two-month period, and we find that these temporary incentives created long-term changes in savings behavior for a substantial portion of those induced to open accounts by the lottery incentives; similarly, Schaner (2016) finds a long-run impact of a temporary incentive that encouraged saving (in her case, a 20% interest rate offered for six months).

Taken together, our results suggest that prize-linked savings accounts can encourage the unbanked to open bank accounts, and that accounts in formal financial institutions may be an experience good. Clients induced to open accounts by temporary lottery incentives appear to use the account initially and learn its value; over time, their initially lower savings balances catch up to those of other users. Nevertheless, a minority remain active account users in the long term, suggesting that if savings accounts are an experience good, the benefit of saving in a formal account is higher than anticipated only for some new account-holders.

Figure 1: Details on Experiment

(a) Timeline of Experiment

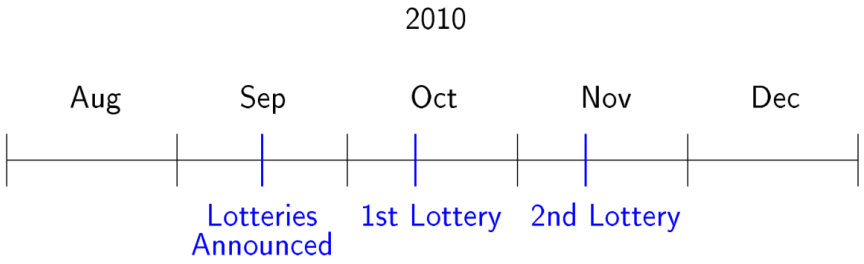

(b) Example Advertisement

Figure 2: Treatment and Control Localities

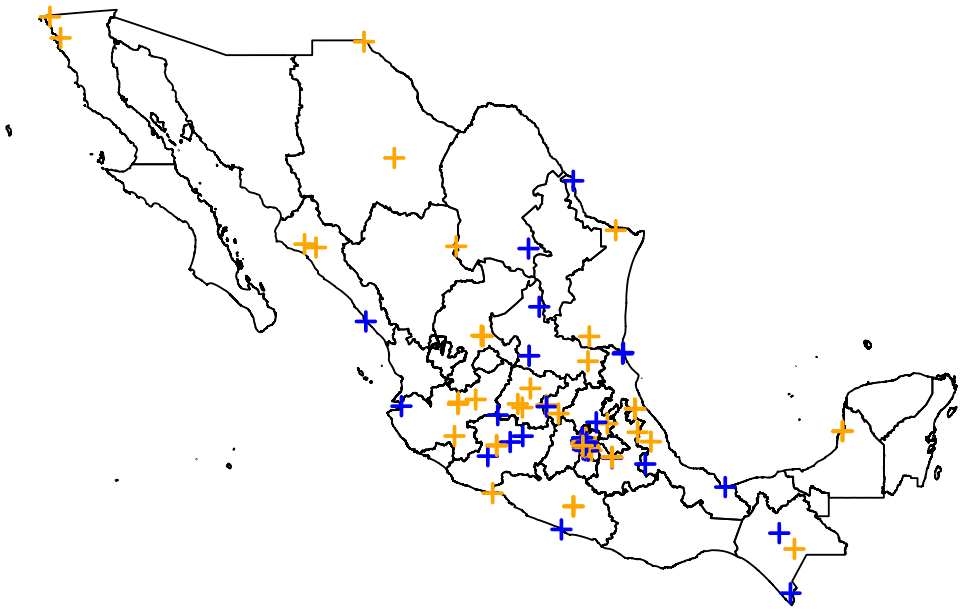

Figure 3: Impact of Treatment on Number of Accounts Opened

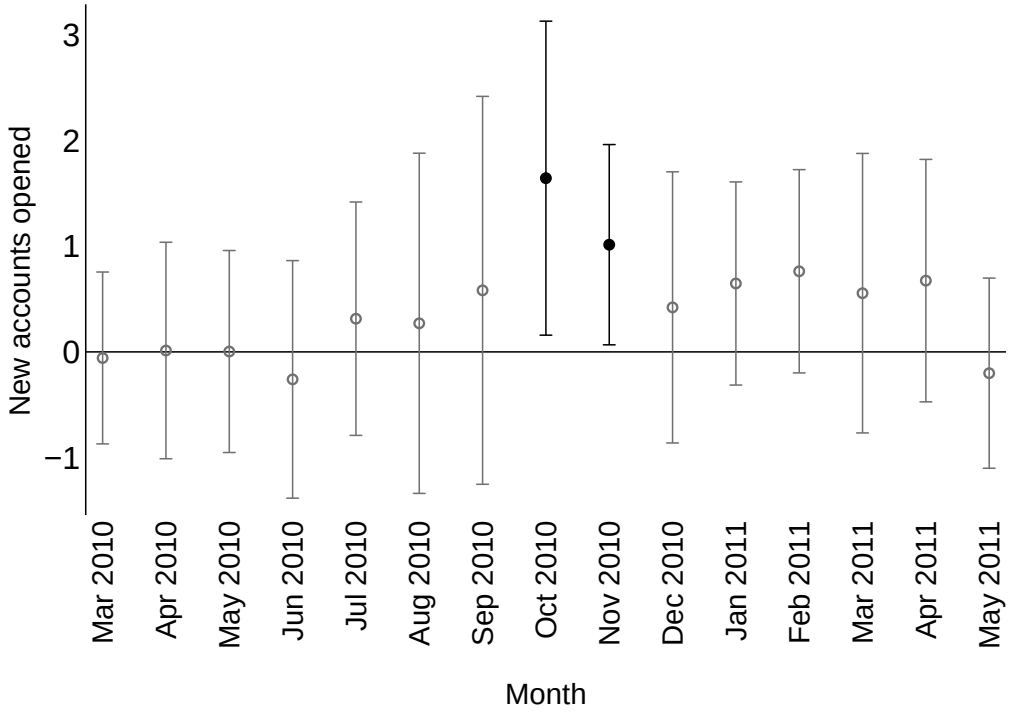

Figure 4: Savings in Existing Accounts in Treatment vs. Control Branches

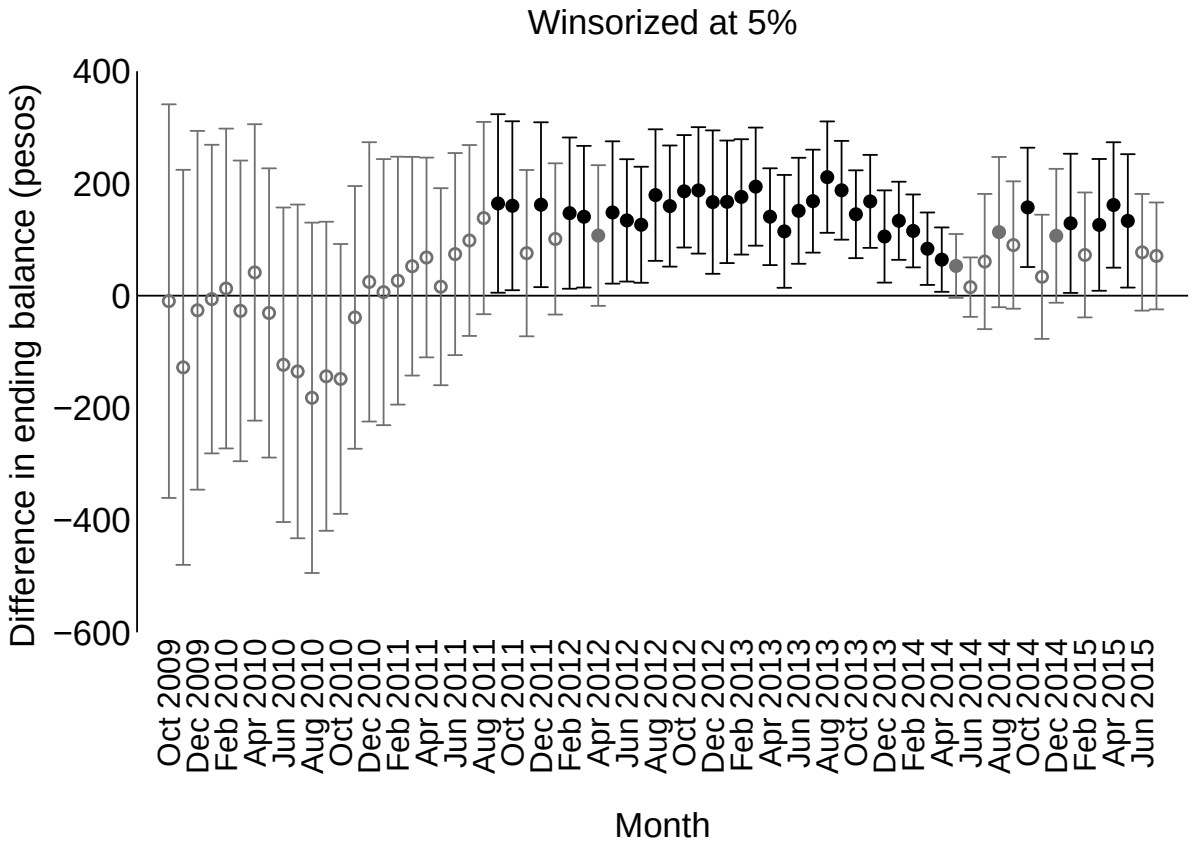

Figure 5: Savings in Accounts Opened October 2010 in Treatment vs. Control Branches

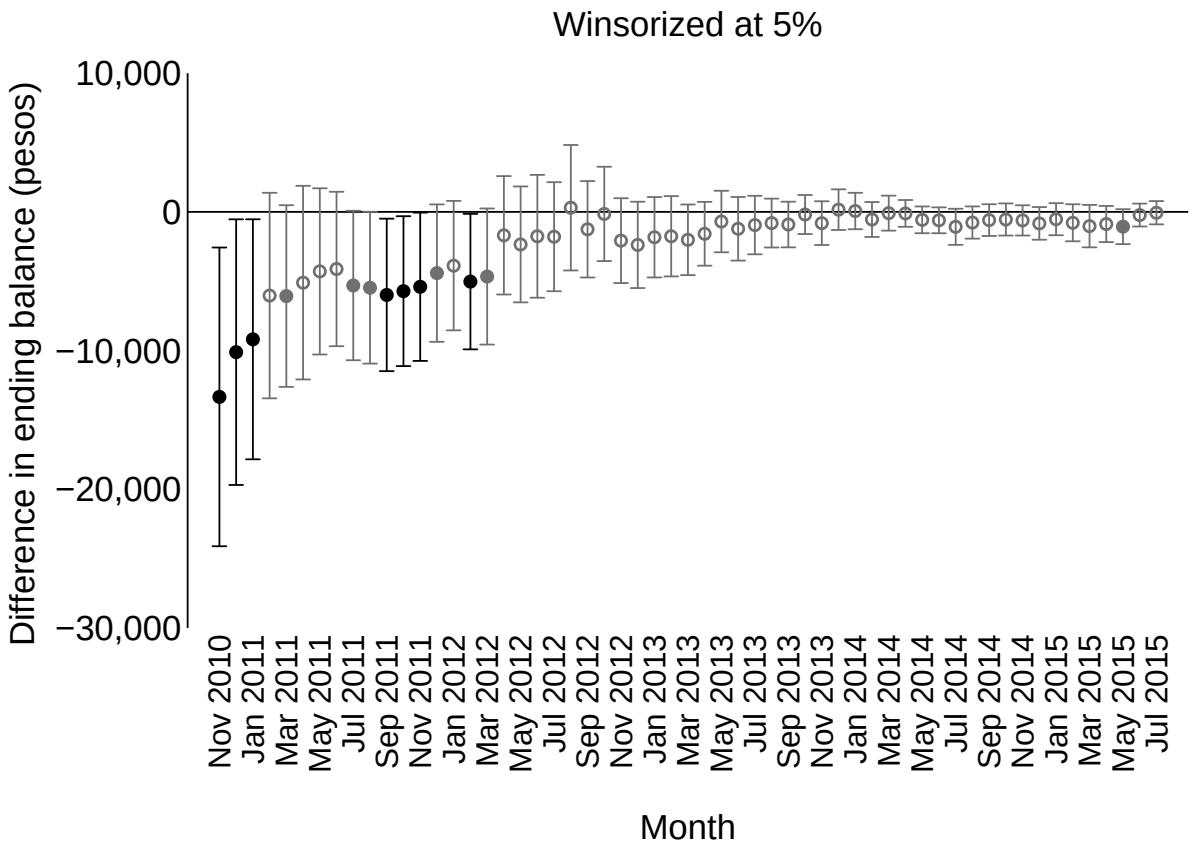

Figure 6: Proportion of Accounts Opened October 2010 Remaining Used in Treatment vs. Control Branches

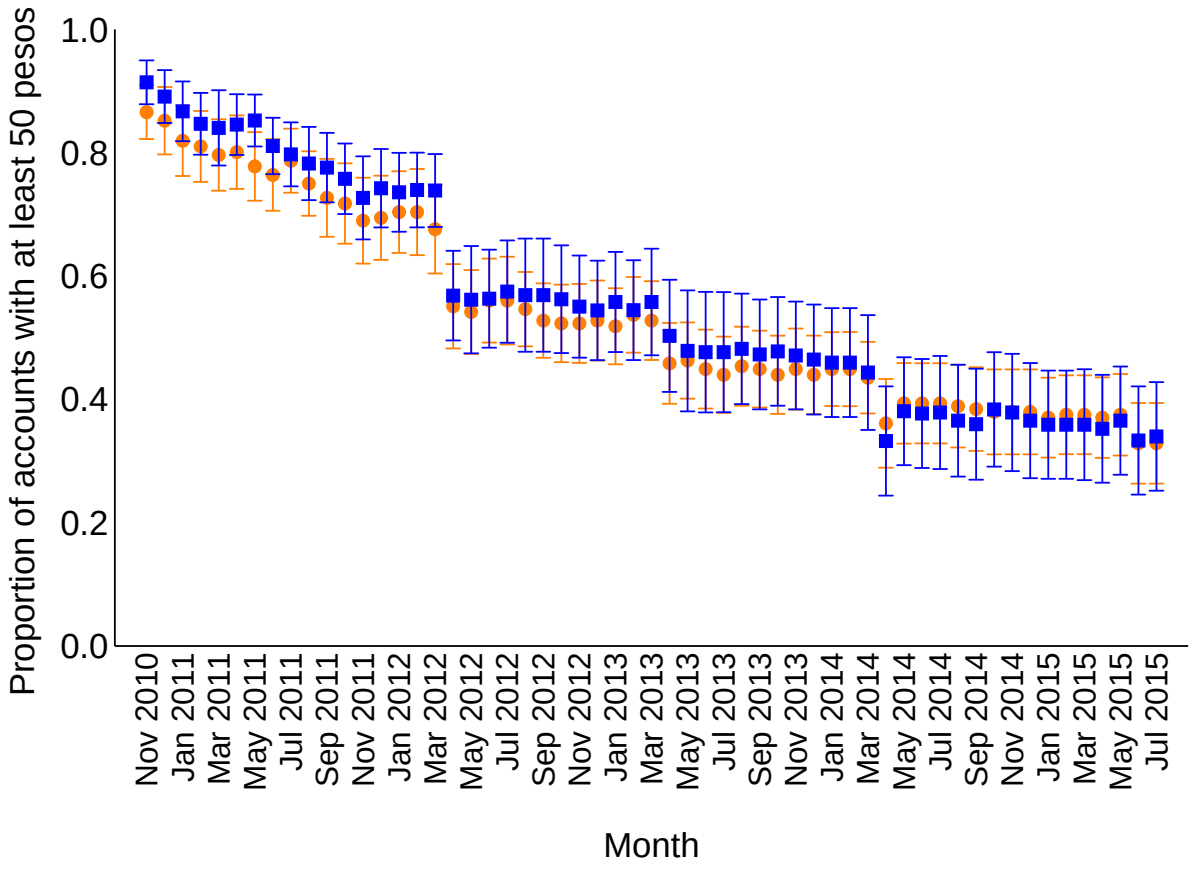

Figure 7: Active Users, Accounts Opened October 2010

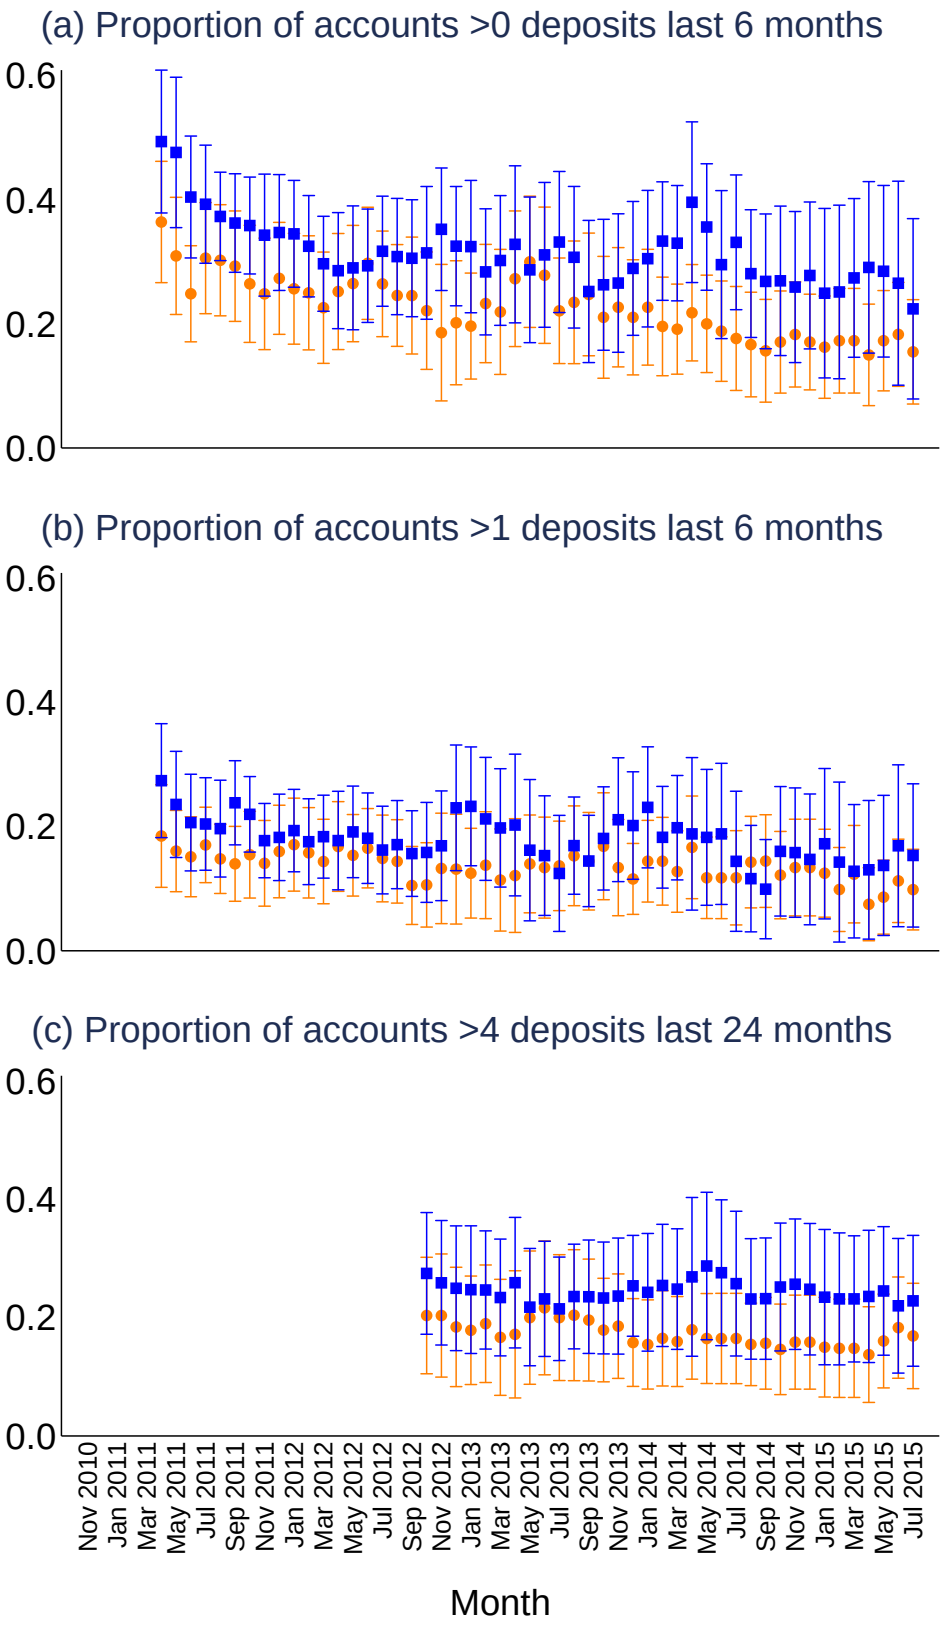

Figure 8: Distribution of Transactions, Accounts Opened October 2010

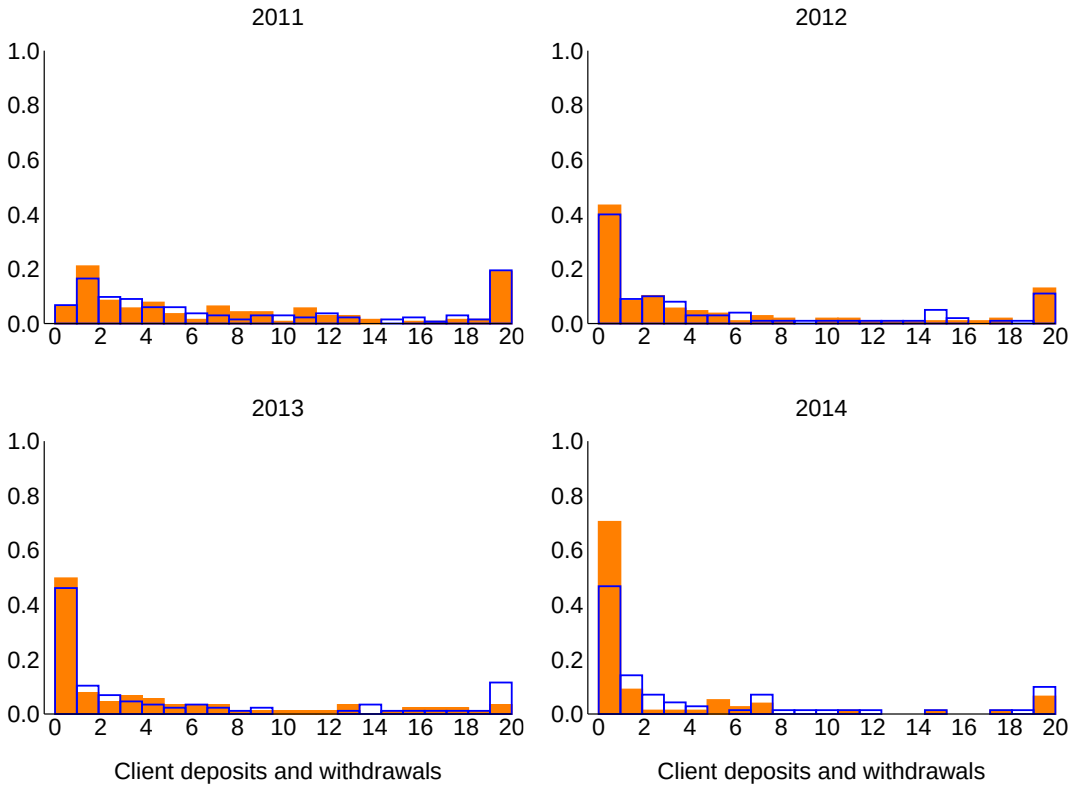

References

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

Atalay, K., Bakhtiar, F., Cheung, S., Slonim, R., 2014. Savings and prize-linked savings accounts. Journal of Economic Behavior & Organization 107, 86–106.

Bachas, P., Gertler, P., Higgins, S., Seira, E., 2017. Banking on trust: How debit cards enable the poor to save more. NBER Working Paper 23252.

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

Burgess, R., Pande, R., 2005. Do rural banks matter? Evidence from the Indian social banking experiment. American Economic Review 95, 780–795.

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

Cole, S., Iverson, B., Tufano, P., 2016. Can gambling increase savings? Empirical evidence on prize-linked savings accounts. Working paper.

Cole, S.A., Tufano, P., Schneider, D., Collins, D., 2007. First National Bank's golden opportunity. Harvard Business School Case 208-072.

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

Demirgüç-Kunt, A., Klapper, L., Singer, D., Oudheusden, P.V., 2015. The Global Findex Database 2014: Measuring financial inclusion around the world. Policy Research Working Ppaer 7255.

Dizon, F., Lybbert, T.J., 2017. Leveraging the lottery for financial inclusion: Lotto-linked savings accounts in Haiti .

Dupas, P., Green, S., Keats, A., Robinson, J., 2016. Challenges in banking the rural poor: Evidence from Kenya's Western Province, in: Edwards, S., Johnson, S., Weil, D.N. (Eds.), Modernization and Development. University of Chicago Press, Chicago.

Dupas, P., Karlan, D., Robinson, J., Ubfal, D., forthcoming. Banking the unbanked? Evidence from three countries. American Economic Journal: Applied Economics .

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

Filiz-Ozbay, E., Guryan, J., Hyndman, K., Kearney, M., Ozbay, E.Y., 2015. Do lottery payments induce savings behavior? Evidence from the lab. Journal of Public Economics , 1–24.

Herskowitz, S., 2016. Gambling, saving, and lumpy expenditures: Sports betting in Uganda. Working Paper.

Jiménez Bautista, S., 2012. Concentración bancaria en México, 1991-2001. Debate Económico 1, 93–117.

Karlan, D., Ratan, A.L., Zinman, J., 2014. Savings by and for the poor: A research review and agenda. Review of Income and Wealth 60, 36–78.

Karlan, D., Zinman, J., 2016. Price and control elasticities of demand for savings. Working Paper. URL: http://karlan.yale.edu/sites/default/files/savingselasticities_dec2016_0.pdf.

Kast, F., Meier, S., Pomeranz, D., 2016. Saving more in groups: Field experimental evidence from Chile. Harvard Business School Working Paper 12-060. URL: http://www.hbs.edu/faculty/Publication%20Files/12-060_8c16f5e7-6fa1-48cc-858d-bba5f12c28ba.pdf.

Kast, F., Pomeranz, D., 2014. Saving more to borrow less: Experimental evidence

from access to formal savings accounts in Chile. Harvard Business School Working Paper 14-001. URL: http://www.hbs.edu/faculty/Publication%20Files/14-001_2237a8d2-6147-4667-9396-cd5ae9702f03.pdf.

Kearney, M., Tufano, P., Guryan, J., Hurst, E., 2011. Making Savers Winners: An Overview of Prize-Linked Saving Products. Oxford University Press, United Kingdom. doi:10.1093/acprof:oso/9780199696819.003.0012.

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

Schaner, S., 2016. The persistent power of behavioral change: Long-run impacts of temporary savings subsidies for the poor. NBER Working Paper 22534.

Schaner, S., forthcoming. The cost of convenience? Transaction costs, bargaining, and savings account use in Kenya. Journal of Human Resources .

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

Tufano, P., Neve, J.E.D., Maynard, N., 2011. U.s. consumer demand for prize-linked savings: New evidence on a new product. Economics Letters 111, 116–118.