

Gambling for Retirement: The Hidden Costs of Savings Lotteries in a Nationwide Experiment*

Matthias Rodemeier

Jared Gars

Justin Holz

Egon Tripodi

Juan Miguel Villa

July 29, 2026

Abstract

Governments increasingly encourage socially desirable behavior with probabilistic incentives. We study such “Pigouvian lotteries” in a retirement-savings field experiment with 387,000 workers in Colombia’s public pension system. The lotteries induce sharp bunching at qualification thresholds, but workers mostly meet them by retiming rather than increasing deposits, leaving overall savings largely unaffected. Lotteries further crowd out valuable life and disability insurance and disproportionately reward wealthier savers. A welfare analysis shows that, once we account for these unintended consequences, lotteries reduce welfare. Our findings illustrate that behavioral spillovers across time and choice domains can reverse the verdict on behavioral public policies.

JEL Codes: D14, D91, H55, C93, G52

Keywords: Retirement Savings, Prize-Linked Savings, Paternalism, Informality, Household Finance

*We are especially grateful to John Beshears, Taha Choukhmane, Stefano DellaVigna, John List, Christopher Palmer and to audiences at the Midwest International Economic Development Conference 2024, NEUDC 2024, ESA Asia-Pacific 2025, SITE 2025, Bocconi, Chicago, Cologne, and Copenhagen for comments that improved the manuscript. Federico Cosenza, Martin Gallardo, and Eleonora Guseletova provided excellent research assistance, and Amanda Slater excellent proofreading. We are grateful to officials from Colpensiones in Colombia for their extensive data sharing and cooperation. The experiments were conducted by Colpensiones. The University of Michigan IRB (HUM00246681) approved the use of secondary data from the lottery experiment, which is registered on the AEA RCT Registry (AEARCTR-0019175). Some members of the research team contributed to the design of the matching experiment, which was approved by the IRB of the University of Essex (ETH2122-1112). Rodemeier: Bocconi University, CESifo, IZA, BAFFI, and IGIER; rodemeier.matthias@gmail.com. Gars: University of Florida and JILAE; jgars87@ufl.edu. Holz: University of Michigan and JILAE; holzj@umich.edu. Tripodi: Hertie School, CESifo and JILAE; tripodi@hertie-school.org. Villa: Inter-American Development Bank; juanmiguelv@iadb.org. Tripodi acknowledges funding by Deutsche Forschungsgemeinschaft through CRC TRR 190 (project number 280092119). The opinions expressed herein do not reflect the opinions of the sponsoring agencies. All remaining errors are our own.

Since Pigou (1920), economists have studied the use of deterministic incentives, such as taxes and subsidies, to optimize resource allocation. One alternative to these policies is the use of *probabilistic* incentives that reward socially desirable behavior with the chance to win a prize. These “Pigouvian lotteries” have gained widespread adoption in many domains, such as vaccinations, energy savings, public transport use, tax compliance, and savings.¹ And they have practical appeal for policymakers: the prize is fixed in advance, so the fiscal cost of the incentive does not grow with the number of people who respond, and a small chance of a large prize may be more attractive to savers than deterministic incentives or actuarially equivalent interest payments (Filiz-Ozbay et al., 2015; Kearney et al., 2010).

Whether such lotteries deliver on this appeal, especially in a high-stakes setting like saving for retirement, turns on questions for which causal evidence is limited: Do lotteries generate new savings or merely reshuffle the timing of deposits? Do they interfere with adjacent retirement saving policies? Do they target those workers whose undersaving motivates intervention in the first place? And how persistent are their effects?

We study these questions in a lottery intervention that was rolled out in 2019 as a nationwide natural field experiment by the Colombian public pension fund, *Colpensiones*, and we observe workers in this intervention for a six-year period. The pension fund invests contributions made by informal workers into a conservative inflation-protected portfolio. The main economic frictions that these savings vehicles try to address are household myopia regarding future retirement needs, the fiscal externality from entering retirement without adequate savings, which results in publicly funded support, and, particularly in economies with high informality, the scarce attention workers can devote to retirement saving, which, in the absence of automatic payroll deductions, becomes a frequent and deliberate act (Beshears et al., 2018; Gomes et al., 2021).²

In the main experiment, the pension fund randomly assigned 387,449 workers to different lottery treatments, offering them the chance to win prizes if they deposited a certain amount in their illiquid retirement account. Colpensiones randomly offered workers one of three prize lotteries via mobile phone text messages, and used no other channels to communicate about the lotteries, so as to maintain random exposure to the treatments and avoid contamination. The *Cash* lottery treatment

¹For example, Sweden rewarded law-abiding drivers with entries into a cash prize draw funded by speeders’ fines (Schultz, 2010), Palo Alto’s energy-savings raffle used lotteries to encourage energy conservation (Browne, 2016), and Indian Railways’ Lucky Yatra campaign turned train ticket purchases into lottery tickets (Team Impact, 2025). Other examples of such lotteries include Norway’s recycling lottery (Norsk Tipping, 2021), Ohio’s “Vax-a-Million” vaccine lottery (Walkey et al., 2021), Singapore’s INSINC transit lottery (Moleskis, 2023), and Hong Kong’s vaccine lottery (Agence France-Presse, 2021). For examples of savings lotteries see, e.g., Guillén and Tschoegl (2002). Savings lotteries also became legal in the United States in 2014 under the American Savings Promotion Act (United States Congress, 2014).

²Designing pension programs and incentives to increase participation is a major policy issue in Colombia and similar countries that face rapid demographic aging, widespread informality, and low pension coverage. Workers entering retirement without adequate savings create fiscal externalities through publicly funded support program enrollment.

offered workers the chance to win a small cash prize if they deposited at least 55,000 Colombian pesos (COP; 17 USD, or about 15% of biweekly income) within a two-week period.³ These workers received 18 consecutive lottery offers over the study period, one every two weeks. The *House* lottery treatment offered workers the chance to win a house if they deposited at least 500,000 COP (152 USD, or about 12% of six-month income) within a six-month qualification period. These workers received two lottery offers during the study, one per semester, though they were messaged on the same biweekly schedule as the other arms. The *House+Cash* treatment offered workers both the biweekly cash lottery and the semiannual house lottery, with the same thresholds and timelines as above. Finally, the *Control* group received biweekly reminders to save but no lottery offers.

We first document that lotteries change savings behavior. The *Cash* treatment significantly increases the probability of savings above the cash lottery qualification threshold by 10.3%. Similarly, the *House* treatment increases the probability of savings above its qualification threshold by 15.7%. Locally, effects of the lotteries are even stronger: the probability that savings bunch *exactly* at the qualification threshold (i.e., at 55,000 and 500,000 COP) increases by 446% for the *Cash* treatment and 71.7% for the *House* treatment; combining both lotteries (*House+Cash*) has intermediate effects. At first sight, the policy thus appears to have its intended effect. The savings dynamics, however, tell a different story: the additional deposits made during qualification windows are offset by missing deposits after the lottery deadline has passed. Workers shift the *timing* of their deposits rather than meaningfully increase their net savings. Accounting for these intertemporal spillovers, the overall effect on savings in the year of the intervention is only 2.3% for *Cash* and 2.0% for the *House* group.

While these intertemporal spillovers render the immediate effects on savings small, temporary lottery incentives may still be valuable if they establish lasting saving habits. Following workers for four years after the intervention, we find no evidence of such persistence: treated subjects save no more than the control group.⁴

The lotteries also interfere with existing retirement policies, likely by crowding out attention. As a hybrid social security system, BEPS offers informal workers free life and disability insurance if they make at least six deposits throughout the year or reach a savings threshold. The lotteries cause workers to consolidate deposits around the lottery qualification thresholds, which reduces their likelihood of meeting the six-deposit requirement for insurance coverage. This behavior is likely a mistake because workers could simply deposit more frequently during the lottery period,

³We report amounts in Colombian pesos (COP) and convert to USD at the Banco de la República annual average exchange rate of 3,281 COP/USD in 2019.

⁴We further do not find any evidence that treated workers are more responsive to future (non-randomized) lotteries that Colpensiones used during the post-intervention period. This does not support the hypothesis that savings lotteries cause addictive behavior in the spirit of [Becker and Murphy \(1988\)](#)—a natural concern with policies that use gambling to change behavior. A caveat is that we cannot rule out whether savings lotteries affect demand for other lotteries we do not observe.

thereby qualifying for *both* the lottery and the life insurance. This crowd-out is a second spillover, now across policy domains: an incentive aimed at savings erodes coverage in a separate insurance program, and survey and revealed-preference evidence show the resulting welfare loss is large, exceeding the lottery’s benefits. Because our administrative data link the two, we can quantify a cost that a savings-only evaluation would miss entirely.

To understand how well-targeted lotteries are, we obtain data on stated time preferences, as well as administrative data proxies for wealth. Results suggest that impatient workers have no larger treatment effects than patient ones. In addition, lotteries tend to be regressive as wealthier workers are more likely to qualify for both the cash and house lotteries. This reflects how the lotteries were designed: rather than rewarding *changes* in savings, they reward *levels* of savings, which favors wealthier workers who already save more.

We close with a calibrated welfare calculation for the two lottery interventions. We recover workers’ revealed savings preferences from a second, complementary field experiment, on the same population, that randomizes deterministic savings matches. We identify the curvature of the utility function from savings responses to different match levels, which allows us to value both the lottery and the insurance. The expected utility from the lottery prize is substantially smaller than the insurance value. Indeed, we find that lotteries reduce social welfare almost entirely due to their crowd-out of life and disability insurance coverage. The welfare computation also uses workers’ subjective winning probabilities as an input, which we recover by comparing savings responses to the stochastic lottery with those to the deterministic match, following [Rodemeier \(2025\)](#) and [Miller, Sahni, and Strulov-Shlain \(2026\)](#). We find that workers overestimate the odds of large prizes (the house lotteries) and underestimate the odds of small prizes (the cash lotteries), consistent with theories in which prizes and probabilities compete for attention ([Barberis, 2013](#); [Bordalo et al., 2012](#); [Kőszegi and Szeidl, 2013](#)), and with the classic observation that rare events are either ignored or overweighted ([Kahneman and Tversky, 1979](#)). Quantitatively, however, these belief distortions matter much less for welfare than the insurance crowd-out. The broader implication of our welfare analysis is that evaluations of behavioral policies must account for their spillovers, over time and across policy domains, not only their direct effect on the targeted behavior.

Our paper contributes to the literature on lotteries as instruments to promote savings and, more generally, socially desirable behavior. Lottery incentives have been used in a wide range of policy settings, including charitable giving ([Landry et al., 2006](#); [Lange et al., 2007a,b](#)), tax compliance ([Naritomi, 2019](#)), and vaccinations ([Campos-Mercade et al., 2024](#)). In the savings domain, prior work has studied government lottery bonds ([Tufano, 2008](#)) and the “prize-linked savings” accounts offered by banks and credit unions ([Cole et al., 2022](#); [Cookson, 2018](#); [Gertler et al., 2023](#)). Yet field-experimental evidence is scarce.⁵ [Gertler et al. \(2023\)](#) find that a temporary deposit lottery at

⁵Important work by [Atalay et al. \(2014\)](#) and [Filiz-Ozbay et al. \(2015\)](#) studies savings lotteries in the laboratory and

a commercial bank in Mexico attracts new accounts, mostly from previously unbanked households, whose deposits persist beyond the incentive; [Bauer et al. \(2022\)](#) show that lottery incentives induce pension-fund participants to view their pension information, but without downstream effects on savings; and [Burke \(2021\)](#) finds that prize-linked incentives for debt repayment attract borrowers who were already likely to repay, with no average effect on debt. We complement this literature by evaluating savings lotteries as *public policy*. Beyond their effects on saving, we study interactions with complementary policies, distributional consequences, and welfare implications. Taken together, this offers a comprehensive normative assessment of savings lotteries as a tool for public finance.⁶

Our paper also contributes to the literature on the long-run consequences of temporary incentives. A common promise of behavioral incentives is that a temporary intervention can durably change behavior ([Gneezy et al., 2011](#)). Temporary incentives have been found to build lasting habits, including in exercise ([Charness and Gneezy, 2009](#)) and handwashing ([Hussam et al., 2022](#)), to durably increase vaccination ([Campos-Mercade et al., 2024](#)) and health-product adoption ([Dupas, 2014](#)), and to generate persistent gains in entrepreneurship and income ([Schaner, 2018](#)).⁷ Our setting offers an unusually long window on this question: we observe every account for four years after the last randomized lottery. We find that the lotteries reorganize the timing of deposits within the incentivized year and leave no detectable effect on savings afterward, and the only effect that does show some persistence is the disruption in the schedule by which workers save.

Finally, our work relates to a literature on the unintended consequences of public policy due to behavioral frictions. Interventions that succeed on their targeted margin can backfire on adjacent ones: tax subsidies for retirement saving are largely offset by reshuffling across accounts ([Chetty et al., 2014](#)), higher minimum contributions to retirement accounts reduce savings held elsewhere and increase borrowing ([Choukhmane and Palmer, 2024](#)), and reminders that reduce credit-card late fees increase overdraft fees on users' checking accounts ([Medina, 2021](#)).⁸ A growing behavioral literature identifies scarce attention as a mechanism behind many such spillovers: cognitive resources are scarce, and people allocate them imperfectly (e.g., [Hanna et al., 2014](#); [Rodemeier and Löschel, 2025](#)). In laboratory experiments, [Altmann et al. \(2022\)](#) show that steering attention to one decision improves choices in that domain but deteriorates them in others, leaving overall

finds substantially larger effects on savings than ours.

⁶[Lockwood et al. \(2024\)](#) evaluate state-run lotteries used to raise tax revenue. This is conceptually different from the corrective (“Pigouvian”) lotteries we study, whose purpose is to incentivize behavior that governments consider socially desirable. Relatedly, [Kuhn et al. \(2011\)](#) investigate how state-run lotteries affect behavior of the winner and the social circle around them.

⁷Yet such effects often prove fragile: exercise habits largely dissipate unless supported by commitment devices ([Royer et al., 2015](#)), and default-induced saving does not persist once payments switch back to cash or workers change employer ([Choukhmane, 2025](#); [Somville and Vandewalle, 2018](#)).

⁸Whether such cross-account effects arise is an empirical question, as in fact [Medina and Pagel \(2025\)](#) find that savings reminders raise savings without increasing borrowing on other accounts.

payoffs essentially unchanged. In our setting, lotteries concentrate attention on meeting qualification thresholds, and in doing so draw it away from an insurance program that rewards deposit *frequency*. A fully attentive worker would have qualified for both by splitting the same total savings across more deposits, thereby preserving insurance eligibility while still reaching the lottery threshold. Our welfare analysis shows that a policy that may look innocuous on its targeted margin is welfare-reducing once these attentional spillovers are priced.

1 Experimental Setting

1.1 Partner Institution

The Colombian government launched *Beneficios Económicos Periódicos* (BEPS) in 2015 to provide a voluntary retirement savings program for informal workers, who make up 60% of the Colombian workforce (OECD, 2024). Workers' contributions into BEPS accounts are invested in an inflation-protected, fixed-income portfolio. Once they reach the minimum retirement age, workers are eligible to receive annuities paid every two months, in amounts based on the total portfolio value and a one-time government match of 20% on their cumulative contributions (Villa and Benitez, 2023).⁹ BEPS accounts are managed by *Colpensiones*, Colombia's state-owned pension administrator.

As a hybrid social security scheme, BEPS also provides access to free basic life and disability insurance. To qualify, account holders are required to either (i) make at least six deposits per year or (ii) reach the target minimum yearly contribution. Because the insurance rewards the frequency of deposits, while the lotteries only reward reaching a specific savings threshold, the two incentives are not mutually exclusive and need not compete with one another—an important detail we return to in Section 2. The target is equal to six days of a legal monthly minimum salary, which was 157,000 COP in 2018 and 166,000 COP in 2019. The insurance covers physical disability, chronic illness, and death, and includes a grant for funeral services for a full calendar year, with an average claims payout of 5 million COP across 2016–2022 (1,500 USD). Coverage applies for a given calendar year, conditional on meeting one of the two eligibility conditions in the prior year.

Fiscal externalities The fiscal externality of undersaving in Colombia arises because workers who approach the minimum retirement age (three years prior) without qualifying for a contributory pension or without having accumulated enough savings in the BEPS system become eligible for the non-contributory old-age program *Colombia Mayor*. This program provides a social assistance transfer of about 80,000–100,000 COP per month, which corresponds to around 60% of the extreme poverty line. Although modest, these transfers represent a fiscal burden since they are funded from

⁹The minimum retirement age is 57 for women and 62 for men. Workers can claim their BEPS annuity from these ages onward, and actual exit from work and claiming typically occur later.

general revenues, and they currently reach close to two million people. Thus, every peso of private undersaving increases the likelihood of public support through *Colombia Mayor*, generating a fiscal externality that justifies Pigouvian incentives to save, above and beyond concerns about household myopia and present bias.

1.2 The Lottery Experiment

To encourage savings by account holders, Colpensiones implemented a savings lottery program in April of 2019. The program involved two types of lotteries: (i) small biweekly “cash prizes” and (ii) large “house prizes” twice a year. The experiment randomized the near-universe of BEPS savers (387,449 account holders as of end-2018)¹⁰ into four groups: three treatments introducing lotteries—a *Cash* treatment, a *House* treatment, and a *House+Cash* treatment—and a *Control* group. Subjects were informed about the respective lottery via biweekly text messages on their phones. All treatment messages included a general reminder to save. To account for the effect of the reminder on savings and isolate the lottery effect, Colpensiones also sent a general reminder to Control group subjects.¹¹

Cash lotteries A total of eighteen cash lotteries took place every two weeks on Wednesdays between April 24, 2019 and December 18, 2019. Every lottery draw consisted of 200 cash prizes, each with a value of 100,000 COP (30 USD). To be eligible for a lottery ticket, account holders were required to save 55,000 COP (17 USD) over two weeks, ending on the Sunday prior to the lottery. The first such window opened on April 8.¹² They could earn a ticket for every 55,000 COP saved, which corresponds to 7.7% of the average monthly informal wage. There was no cap on the number of tickets each worker could earn.

House lotteries The large house lotteries each consisted of one prize, a housing voucher worth up to 102,000,000 COP (31,100 USD), which could be applied toward the purchase of a house, all associated fees, furniture, and related equipment. The house lotteries took place twice, one per semester, drawn on July 3, 2019 and January 15, 2020. As with the cash lotteries, workers earned one ticket for every 500,000 COP (152 USD, 70% of the average monthly informal wage) saved, with no cap on the number of tickets they could earn. Qualifying deposits accumulated over each semester: January 1 to June 30 (July 1 to December 31) for the first (second) lottery. Subjects were

¹⁰BEPS was still a young program at the time but has since grown to over two million enrollees as of late 2025.

¹¹Reminders are themselves an active ingredient of this program: as we saw in a pilot study, personalized savings reminders significantly increase the frequency and amount of BEPS deposits (Gars et al., 2025).

¹²The two-week qualification window opened April 8 (three days before messaging began on April 11) and closed Sunday, April 21, three days before the April 24 draw.

informed of the first house lottery only in April, about three months before its deadline,¹³ but their cumulative savings since January 1 counted toward the threshold; for the second lottery they were informed at the start of the qualification period. Both lotteries therefore use six months of savings to determine whether subjects qualified.

The 387,449 accounts were split roughly evenly across the four arms (see Appendix Table A1).¹⁴ Since the results of the *House+Cash* group turn out to be a combination of the two separate lotteries, we focus on the individual lotteries in the main part of the paper and defer results for *House+Cash* to the Appendix. Beginning on April 11, 2019, all four arms received SMS messages on the same biweekly schedule. Every message included a general reminder to save and made no mention of the life and disability insurance (see Table A2 for an English translation of these messages). Workers received no notification of qualification or of the tickets they had earned; under the lotteries' official terms and conditions, only winners were contacted. Subjects were not told how many lotteries there would be or when the program would end, but over time they likely inferred the offer would recur.

In terms of technical implementation, Colpensiones randomized the text messages rather than the lotteries, meaning that each draw included all workers who qualified for that lottery, independent of their assigned treatment. However, because the text messages were the only way Colpensiones informed workers about a lottery offer, this randomization is essentially as if the lotteries were directly randomized. The main threat to this interpretation is the possibility of spillovers, whereby workers assigned to one lottery may learn about another from other subjects. Similarly, workers in the *Control* may learn about the lottery from friends in the treatment groups. Under the plausible restriction that spillovers weakly increase lottery exposure among *Control* workers, any such contamination attenuates the estimated treatment effect toward zero, implying that our estimates provide a lower bound on the causal effect of lottery exposure (Vazquez-Bare, 2023). That said, in Section 2, we provide empirical evidence *against* treatment spillovers. Savings of account holders in each treatment group bunch at the qualification threshold of the randomly assigned lottery treatment and *not* at the threshold of the other lottery.

Post-experimental lottery programs In the years after the experiment, savings lotteries have become a recurring component of BEPS's outreach strategy. Each year, BEPS launched additional promotional lotteries and advertised them widely through billboards, television programming, and social media channels. Figure 1 provides examples of these advertisements.¹⁵ Crucially, these post-2019 promotions were not randomized: they were offered to all account holders, so our long-

¹³This is simply because Colpensiones decided to implement the policy after the first quarter of the year but wanted to split the two lotteries into semesters.

¹⁴Due to administrative capacity, a random subset of *Cash* and *House+Cash* accounts received lottery messages only during the first three months (April–June) while the remainder received messages until the end of 2019. Our analysis focuses on the full-dose groups unless otherwise indicated.

¹⁵Post-experimental lotteries differed in prize size but always had a one-year deadline, at the end of each year.

Figure 1: Pictures from BEPS Savings Lotteries



(A) One of the House Lottery Winners in the 2019 Experiment



(B) Advertisement Billboard in a Follow-up Year



(C) YouTube Video in 2024

Notes: This figure shows pictures from BEPS savings lotteries. Panel A displays a house lottery winner at the end of our experiment in 2019. Panels B and C show non-randomized advertisements used in subsequent years, illustrating how BEPS promotes savings lotteries nationwide through billboards, TV shows, and digital media.

run analysis (Section 2) will capture how variation in exposure in 2019 affects subsequent saving behavior when lottery incentives become universal. These large-scale policy campaigns, which continue to this day, underscore the continued policy importance of understanding how savings lotteries affect behavior and welfare.

Data BEPS provides us with a rich dataset that details each deposit that subjects made to their retirement savings accounts. Every deposit includes a time stamp, the amount deposited, and an account ID. We observe deposits for six years, starting in January 2018 and ending in January 2024. Each BEPS account belongs to a single informal worker, our unit of analysis.¹⁶ Appendix Table A1 reports summary statistics across treatment groups. All observables are balanced across groups.

In total, we observe 387,449 account holders in our data. In the year before the experiment, 2018, the average account holder saved 97,594 COP ($\approx$ 30 USD). Around 68% of account holders saved at least once a year and 70% are female, which is in line with women disproportionately working informally. The average age is 50 years, which is only 7 years below the minimum retirement age for women and 12 years below the minimum retirement age for men. Since BEPS was introduced in 2015, the average account is only around 2 years old. We also observe each account holder’s welfare index, called SISBEN, which is assigned by the Colombian government to identify beneficiaries of social assistance programs. The SISBEN is composed of information on income, health, and access to basic services. It ranges from 0 to 100 (from poorest to richest). The average SISBEN score of account holders in our sample is 31.5, which is considered to capture “moderate poverty.” Since the index is designed to predict the permanent component of household income for targeting purposes (Beuermann et al., 2025), we use it as a proxy for a worker’s economic position or wealth.

2 Results

In this section we document two ways the lottery reshapes savings, over time and across policy domains, and we examine the key rationales for savings lotteries.


The lotteries are a nonlinear incentive that rewards savings above a threshold $\underline{s}$. If subjects respond to this incentive, we would expect an increase in the density of the savings distribution at $\underline{s}$.¹⁷ We test for this pattern by regressing savings indicators on treatment:

$$d_i = \alpha + \mathbf{Z}_i' \boldsymbol{\beta}_Z + \mathbf{X}_i' \boldsymbol{\theta} + \lambda_r + \epsilon_i \quad (1)$$

¹⁶The terms worker, account holder, and subject are used interchangeably across sections. We call an account active if it made at least one deposit in the six months before randomization, and inactive otherwise.

¹⁷In a frictionless world, a rational model would only predict an increase at and not above $\underline{s}$. In a world with frictions and lumpy savings, we also expect workers to choose more diffuse savings above $\underline{s}$. See, for instance, Gennaioli et al. (2026) for formal arguments.

Figure 2: Saving at and above the Qualification Threshold



Notes: This figure presents the effects of the *Cash* and *House* treatments on the likelihood of saving at and above the lottery qualification thresholds, based on equation (1). Panel A plots estimates of the percentage point change in the likelihood of saving exactly (1) 55,000 COP in any of the eighteen cash lotteries and (2) 500,000 COP in either of the two house lotteries. Panel B plots estimates of the percentage point change in the likelihood of saving at least (1) 55,000 COP in any of the eighteen cash lotteries and (2) 500,000 COP in either of the two house lotteries. Each regression controls for baseline savings, age, gender, and account age, with department fixed effects. Error bars denote 95% confidence intervals; standard errors are robust.

where d_i is a binary outcome for account i during the experiment, $\mathbf{Z}_i$ is a vector of treatment indicators, and $\mathbf{X}_i$ is a vector of account-level controls: baseline savings in the year prior to the experiment (total deposits over 2018 and over the pre-intervention months of 2019), age, gender, and account age. We include department fixed effects λ_r and estimate robust standard errors. In our first specification, $d_i = \mathbf{1}\{s_{it} = \underline{s}\}$ and in our second specification we use $d_i = \mathbf{1}\{s_{it} \geq \underline{s}\}$ for any t during the experiment. We run this regression for the cash lotteries where t indicates each of the 18 two-week qualification periods with $\underline{s} = 55,000$ COP, and then separately for the house lotteries where t indicates each of the two six-month qualification periods with $\underline{s} = 500,000$ COP.

2.1 Incentive Effects at the Threshold

Figure 2 presents treatment-effect estimates from equation (1). Panel A displays differences in the probability of saving exactly at the threshold in any cash lottery qualification period. During the experiment, the probability of choosing savings equal to any cash lottery threshold increases by 1.3 percentage points for the *Cash* treatment, which represents a large increase of 446% relative to the *Control* group. Consistent with the absence of treatment spillovers, there is no evidence that subjects in the *House* treatment targeted the cash lottery.¹⁸

¹⁸We also plot the results of estimating equation (1) for each two-week period in Figure A1. We find that accounts in the *Cash* treatment respond immediately to the cash lottery offers by saving at the threshold and the effects decline

Similarly, the *House* treatment increases the probability of six-month savings bunching at the higher savings requirement of 500,000 COP by 0.46 percentage points, which is an increase of 71.7% relative to *Control* accounts. Subjects in the *Cash* treatment, who should not know about the house lottery, indeed do not increase the likelihood of saving at the house lottery threshold.

In Panel B, we look at differences in the probability of qualifying for the lotteries in any of the qualification periods. The impact of the *Cash* treatment on the probability of meeting one of the cash lottery thresholds is around 0.84 percentage points, or 10.3% relative to *Control*. For the house lottery, we find an increase of 0.56 percentage points ($p < 0.001$), which is a 15.7% increase relative to *Control* accounts. We further see that even subjects in the *House* treatment are more likely to pass the threshold for any of the cash lotteries. This is because saving for the house lottery makes subjects mechanically more likely to qualify for a cash lottery. The effect differs importantly from Panel A, where only accounts in the *Cash* treatment had an incentive to save *exactly* at the cash threshold, and only *House* treatment subjects had an incentive to save at the higher threshold. Thus, Panel A isolates the deliberate response to each group’s own threshold, while Panel B also captures this mechanical overlap.

Result 1: *Lotteries cause savings to bunch at and above the qualification threshold.*

2.2 Impact on Aggregate Savings

Next, we examine whether the lotteries had their intended effect of increasing overall savings. Figure 3 shows a time series of the treatment effects on *average savings* as a percentage change relative to the *Control*. When constructing savings outcomes, we aggregate the total value of deposits at the biweekly level for both treatments.

Panel A plots estimated treatment effects for *House*. Treatment and *Control* do not differ before April, in line with successful randomization. Savings increase in the first month after receiving the offer. The effect is sizable in relative terms during the first house lottery, as savings increase markedly in May and June. Importantly, after the qualification period for the first house lottery, the treatment effect becomes *negative* and statistically significant in July and August. This increase before the deadline, offset by a decline right afterward, indicates that deposits are brought forward from future periods in order to meet the threshold—an *intertemporal spillover* of the treatment that undermines the policy objective. Over time, savings recover and the treatment effect remains around zero throughout the rest of the year. The second house lottery does not seem to induce any additional savings response.

Panel B plots the treatment effects on savings for the cash lottery. In contrast to the *House* but persist throughout the experiment.

Figure 3: Short-term Treatment Effects on Savings



Notes: This figure presents the effects of the *House* and *Cash* treatments on total biweekly savings in 2019. Thick lines in Panels A and B denote the intervention periods for the house lottery (Panel A), from April 11 to December 31, 2019, and the cash lottery (Panel B), from April 11 to December 15, 2019, respectively. Each coefficient is from a separate regression where the outcome variable is total savings during biweek t , controlling for baseline savings, age, gender, and account age, with department fixed effects. Effects and confidence intervals are expressed as a percentage of the *Control*-group mean. Each point is plotted at the midpoint of the two-week window it summarizes; the first such window opens on April 8, three days before messaging begins on April 11. Through June 30, 2019, Panel B retains the partial-dose *Cash* accounts (messed only April–June 2019); they are excluded from later periods. The dotted vertical line in Panel A marks the deadline of the first house lottery (June 30, 2019). Filled dots denote estimates with a p -value below 0.05. Shaded bands around the point estimates denote 95% confidence intervals; standard errors are robust.

treatment, we do not find treatment effects to be strong in the first months of the intervention. Instead, the effects of the *Cash* treatment gain traction in October through December. One possible reason why treatment effects kick in toward the end of the year is that savings in BEPS accounts are highly seasonal and peak every year in December, which implies that our lottery incentives are also more top of mind for people. Another possibility is that workers target lotteries close to the Christmas holidays.¹⁹

Similar to Panel A, treatment effects become *negative* after the last cash lottery deadline has passed: deposits are shifted forward into the cash qualification periods and are then missing afterward. The negative dip starts immediately after the deadline of the last cash lottery on December

¹⁹It is conceivable that the treatment effects in the *Cash* treatment become larger over time because workers find out about the house lottery that is running in parallel, and try to target the higher threshold. Several patterns in the data suggest that this is not the case. First, in the *Cash* group we do not observe differential bunching at the house threshold or lottery qualification rates relative to the *Control*, including the second qualification period at the end of the year. Second, when we study the *House+Cash* group in the Appendix, we find it is an exact combination of Panels A and B: savings spike before the first house deadline but also before the end of the year. The latter effect must be driven by the *Cash* treatment (Panel B) because it is not observed in the *House* treatment (Panel A). Third, savings in Panel B spike just before the final cash lottery deadline on December 15, not before the final house deadline on December 31. These pieces of evidence suggest that the increase in treatment effects toward the end of the year in Panel B is driven by the final cash lotteries and not by subjects learning about the second house lottery.

15.²⁰

We present treatment effects on aggregate 2019 savings after April 8 in Appendix Table A3. We find that aggregate savings increase by just 2.3% ($p = 0.242$) in the *Cash* group and by 2.0% ($p = 0.047$) in *House* compared to the *Control* mean (69,404 COP, around 21 USD). These modest treatment effects reflect the pattern in Figure 3, in which workers respond to the lotteries by shifting the *timing* of their deposits rather than substantially increasing total savings.²¹ This highlights the importance of measuring behavior beyond the treatment period and provides our first piece of evidence that a narrow focus on targeted behaviors can overstate the policy success.

Result 2: *Workers shift the timing of their deposits into the qualification periods. This intertemporal spillover eliminates much of the policy impact, as evidenced by the small effects on aggregate savings.*

2.3 Long-Run Effects

While the policy did not substantially increase savings, the leading argument for savings lotteries may be a different one: a temporary incentive may instill a *habit* of saving. Under this view, the temporary reward serves as a catalyst—designed not merely to buy a short-term spike in deposits, but to cultivate lasting financial routines. If such habits take root, even a modest initial response can yield substantial cumulative benefits, making an extended evaluation horizon essential to capturing the true value of the policy. We study whether the treatments induced persistence in savings increases using rich administrative data from BEPS that follow accounts for four years post-treatment, zooming out the panel variables to the quarterly level. We plot the time series of treatment effects on savings in Figure 4, starting in January 2018 and following accounts until January 2024. Each panel presents the effects of a different treatment against the *Control* group. In the years after the experiment, BEPS also ran savings lotteries but they were not randomized and had different thresholds and prizes. The qualification periods after 2019 began on January 1 and had a deadline of December 31.²² These follow-up promotions are nowadays an important part of BEPS’s retirement promotion program (recall Figure 1).


Under habit formation, we would expect positive treatment effects to persist after the experiment ended in January 2020. However, the coefficients are mostly negative and statistically insignificant.

²⁰The negative dip is observed even before the deadline of the last house lottery on December 31, confirming that cash lottery subjects did not target the house lottery.

²¹As shown in Panels B and C of the table, the effect is entirely concentrated among *active* accounts that made at least one deposit in the six months prior to the experiment.

²²In 2020, eight house prizes were drawn from this single annual qualification pool and accounts earned one ticket for each 5,000 COP saved. The threshold for each ticket then increased to 10,000 COP and 20,000 COP in 2021 and 2022, respectively.

Figure 4: Long-Run Effects and Responses to Follow-up Lotteries



Notes: This figure presents the effects of the *House* and *Cash* treatments on total quarterly savings from 2018 through 2023. The thick colored vertical lines in Panels A and B denote the intervention periods for the house lottery (Panel A), from April 11 to December 31, 2019, and the cash lottery (Panel B), from April 11 to December 15, 2019, respectively. The other vertical lines indicate the end of the qualification periods for subsequent lotteries between 2020 and 2023. Each coefficient is from a separate regression where the outcome variable is total savings in quarter t , controlling for baseline savings, age, gender, and account age, with department fixed effects. Effects and confidence intervals are expressed as a percentage of the *Control*-group mean. Each point is plotted at the midpoint of the quarter it summarizes. Through June 30, 2019, Panel B retains the partial-dose *Cash* accounts (messed only April–June 2019); they are excluded from later periods. Filled dots denote estimates with a p -value below 0.05. Shaded bands around the point estimates denote 95% confidence intervals; standard errors are robust.

We also do not observe that accounts treated during the experiment respond differently around the deadline of lotteries in future years. This latter finding informs the question of whether savings lotteries could be addictive: if lotteries caused addiction, randomly exposing subjects to a lottery once would cause treated accounts to be more inclined to respond to future lotteries (Becker and Murphy, 1988), which has been observed in the context of gambling with (non-Pigouvian) state lotteries (Guryan and Kearney, 2010). An important caveat to our conclusion is that we do not observe whether the savings lotteries changed gambling behavior outside of the BEPS accounts, such as in casinos. While we cannot exclude this possibility, existing evidence rather suggests the opposite: Cookson (2018) and Cole et al. (2022) find that the introduction of savings accounts with a lottery element *reduces* unrelated gambling.

Result 3: *The savings response shows no detectable persistence. In the four post-intervention years, treated accounts save no more than Control and are no more responsive to subsequent lotteries.*

Taken together, our short- and long-run findings illustrate that lotteries failed to alter the structural savings behavior of informal workers.

2.4 Impacts on Other Benefits of Saving in BEPS

The way workers meet the qualification thresholds of their assigned lottery can spill over to programs that also reward how *frequently* a worker saves, not just how much. As discussed in Section 1, BEPS also offers access to free life and disability insurance to account holders. To qualify, they are required to either (i) make at least six deposits in a calendar year or (ii) save the target minimum yearly contribution equal to six days of a legal monthly minimum salary (166,000 COP in 2019).²³ Coverage is provided for a full calendar year conditional on qualification in the prior year and includes physical disability, chronic illness, and death, as well as a grant to cover funeral services. A savings promotion like a lottery has the potential to expand this important insurance coverage among the working poor. This could be true even if subjects do not increase aggregate savings but simply deposit more frequently to meet the threshold of six deposits per year. However, the evidence from the previous sections suggests that workers responded to the lotteries by consolidating deposits around the qualification thresholds, raising the concern that the lotteries may have inadvertently reduced insurance coverage.


In Figure 5, we explore this question. For ease of exposition, we pool the *House* and *Cash* treatments, as the effects are similar across conditions. We construct a qualification dummy y_{it} that is equal to one if subject i either made at least six deposits in year t or the total value of deposits was equal to or above the qualification threshold in year t . At least one of the two requirements has to be met to receive insurance coverage in the following calendar year. We plot the estimated treatment coefficients for each year in Panel A.

As expected, there is no difference between treatment and *Control* in 2018, before the messages began. During the treatment year, however, the lotteries cause a statistically significant *drop* of 1.4% in the share of accounts that qualify for insurance relative to *Control*, where the average qualification probability is 26.3% ($p = 0.008$). The reduction in qualification is 2.3% ($p = 0.035$) for the *Cash* and 1.3% ($p = 0.020$) for the *House* treatments. The reduction is driven entirely by the deposit channel: treated workers who target the qualification thresholds consolidate their savings into fewer, larger deposits, reducing their annual deposit count below the six-deposit insurance eligibility requirement. Strikingly, the negative effect on insurance qualification in 2020 is larger than during 2019 (3.2% relative to *Control*), but does not persist in the following years. This can be explained either because subjects have preponed savings into 2019, which are then missing in 2020, or because subjects become accustomed to depositing less frequently.

The welfare loss from not qualifying for this insurance is likely substantial. First, according to a study by Colpensiones, the insurance is highly valued by account holders, with 61% of survey respondents reporting it as their *primary* motivation for saving with Colpensiones (Colpensiones,

²³In 2018, the target amount was slightly lower, equivalent to 157,000 COP.

Figure 5: Eligibility for Life and Disability Insurance



Notes: This figure presents the effects of the *House* and *Cash* treatments on eligibility for life and disability insurance. Panel A shows the effects on annual insurance qualification. The reported percentage changes are based on separate estimates from equation (1), where the outcome equals one if account i qualifies for life and disability insurance in year t , and zero otherwise. Each regression includes a treatment indicator (for *House* and *Cash* pooled) and controls for an active-saver indicator, baseline savings, age, gender, and account age, with department fixed effects. Shaded bands around the point estimates denote 95% confidence intervals; standard errors are robust. Panel B displays the distribution of the total number of deposits made in 2018 (top-coded at 20). The sample includes all accounts that made at least one deposit in 2018 and were opened prior to January 1, 2018. The dashed vertical line in Panel B marks the deposit-count threshold for one of the two possible routes to qualifying for life and disability insurance (i.e., six deposits).

2023). Second, Panel B of Figure 5 displays the empirical distribution of the total number of deposits in 2018 for accounts that made at least one deposit. The distribution clearly shows that account holders change their behavior in response to the insurance incentives: only 4% and 8% of accounts made five and seven deposits, respectively, while 20% made exactly six deposits to meet the benefits threshold.²⁴ This bunching provides revealed preference evidence that people indeed value the insurance.

Our interpretation is that the observed reduction in insurance qualification due to the lottery is likely a mistake rather than a conscious choice by subjects. This is because qualifying for a lottery and qualifying for the insurance are not mutually exclusive, so there is no economic trade-off. For example, instead of making a single lump-sum deposit to hit the lottery threshold, a worker could simply divide that exact same amount into six smaller transactions. By merely increasing the frequency of their deposits—without saving a single peso more—they would successfully qualify for both programs. There is also no meaningful cost of depositing, as deposits can be made digitally. Therefore, the negative effect on insurance qualification suggests that the lottery promotions crowd out attention to other benefits, causing workers to forget about the insurance. This interpretation is consistent with evidence in Rodemeier and Löschel (2025), who show in a different context that

²⁴The other peak at 12 deposits, although much smaller, captures accounts that save monthly.

policies can compete for attention and, therefore, negatively interact.

Result 4: *Lotteries reduce eligibility for life and disability insurance. Since qualifying for the lottery and the insurance are not mutually exclusive, this appears to be a mistake rather than a deliberate trade-off.*

These findings reveal a profound *cross-domain spillover* as concentrating workers’ attention on a narrow savings threshold inadvertently crowds out a highly valuable safety net. Ultimately, this underscores why evaluations of behavioral policies must track choices across both time and interconnected programs to capture their true welfare costs.

2.5 Heterogeneity

While lotteries have no meaningful impact on savings of the *average* worker, they might still increase welfare if they disproportionately target behavior of present-biased or poor workers, whose undersaving a planner most wants to correct (Choukhmane and Palmer, 2024). As gambling has been shown to be more concentrated across these groups (Kearney et al., 2010), savings lotteries could be a well-targeted policy tool. In this section, we study whether lotteries have different effects on these relevant subgroups.

While wealth proxies are directly observed in our administrative data in the form of SISBEN scores, time preferences are not. We follow a common approach in the literature and connect survey proxies of present bias with our administrative savings data (e.g., Lockwood et al., 2024). Prior research finds that these stated preferences predict incentivized choices well (e.g., Falk et al., 2022) and our analysis confirms this insight, as we discuss below. To explore the relationship between the lotteries and time preferences, we use data from the World Bank’s 2012 *STEP Skills Measurement Household Survey*, which provides a representative sample of urban households in Colombia (World Bank, 2014). In addition to questions that measure labor market outcomes and family composition, the survey also elicits stated time preferences. We link survey responses to our administrative data through observable characteristics available in both. Specifically, we limit the World Bank Survey data to informal workers and match observations by age (in five-year bins), gender, and regional administrative division (called “departments” in Colombia).²⁵ This procedure yields imputed measures of stated preferences for 96,516 workers in our experiment.²⁶

²⁵Colombia is made up of 1,104 municipalities organized into 32 departments and one capital district. The 2012 Colombia STEP survey collected a representative sample of adults aged 15 to 64 living in 13 major metropolitan areas that spanned 21 municipalities and 9 districts (World Bank, 2014).

²⁶Since the World Bank survey covers fewer geographic areas than our administrative data and samples workers under the age of 65, we cannot assign survey responses to every worker in the experiment. While this might introduce selection bias, we actually find that whether survey responses are missing does not correlate with treatment assignment.

Measures of myopia The survey uses multiple measures that are potentially informative about myopic behavior and highly predictive of actual savings in the data. The first one is a coarse version of the “staircase method” in experimental economics: respondents first choose between receiving 2,160,000 COP today versus 3,240,000 COP in one year (a 50% return). Those who choose the immediate payment are then asked if they would wait for 4,320,000 COP (100% return), while those who choose to wait are asked if they would still wait for only 2,592,000 COP (20% return). The survey also asks those with discount rates above 100% to state the minimum amount they would accept in one year instead of the immediate payment, which avoids censoring discount rates of very impatient workers. A caveat of the survey is that it only allows us to compute a single discount rate (rather than two separate ones required to point-identify time inconsistency). While any high discount rate could in principle reflect rational exponential discounting, the literature has found that high discount rates strongly predict present-biased behavior (e.g., Toubia et al., 2013).²⁷

In addition, the survey elicits personality traits that have natural connections to self-control problems. The relevant questions are: “*Do you work very hard? For example, do you keep working when others stop to take a break?*”, “*Do you prefer relaxation more than hard work?*” and “*Do you think about how the things you do will affect you in the future?*” Respondents answer on a four-point scale from “*almost always*” to “*almost never*.”

In Table 1, we present the results of a set of regressions that interact treatment assignment with wealth and our survey-based proxies of present bias. Column 1 examines heterogeneity by wealth, while columns 2–5 explore whether treatment effects vary with measures of time inconsistency: the discount rate from the sequential bracketing task and the three personality traits related to self-control and forward-looking behavior. Each column estimates a specification of the form:

$$Y_i = \alpha + \beta_1 \text{House}_i + \beta_2 \text{Cash}_i + \mu H_i + \delta_1 (\text{House}_i \times H_i) + \delta_2 (\text{Cash}_i \times H_i) + \mathbf{X}_i' \boldsymbol{\theta} + \varepsilon_i \quad (2)$$

where Y_i denotes savings as a percentage of the *Control* group mean, H_i is the heterogeneity dimension of interest (an indicator for above-median wealth, impatience, or a personality trait), and $\mathbf{X}_i$ includes controls for baseline savings and, where applicable, for wealth and impatience. When estimating the correlations with time preferences, we always control for wealth in order to avoid

We further do not find that missing survey responses correlate with other observable characteristics.

²⁷Estimating present bias directly requires observing both a near-term tradeoff (e.g., today vs. in one year) and a tradeoff between two future points in time (e.g., in one year vs. in two years). Since we do not observe the latter choice, we interpret high discount rates from the first choice as evidence of present bias. In our main analysis, we classify workers above the median discount rate of 160% as impatient—i.e., those subjects who are only indifferent between the two payments if the later payment is 2.6 times as large as the payment today. While we view this as a conservative definition of present bias, we obtain similar results to the ones discussed in this section when we restrict present bias to subjects with much higher discount rates than the median. Further, questions that ask for a qualitative self-evaluation of personal characteristics (e.g., enjoying relaxation more than hard work) correlate intuitively with high discount rates and yield similar results regarding the targeting properties of lotteries.

Table 1: Heterogeneous Effects of Lotteries on Savings, by Wealth and Myopia

	(1)	(2)	(3)	(4)	(5)
	Wealth	Impatience	Prefers work	Works hard	Thinks of future
<i>House</i>	1.070 (1.053)	4.002** (1.644)	2.075 (1.506)	3.560** (1.537)	3.419** (1.585)
<i>Cash</i>	4.805** (2.141)	-0.120 (2.958)	2.307 (2.836)	2.282 (2.914)	4.626 (3.040)
Wealth	3.824*** (1.165)	5.000*** (1.034)	5.018*** (1.034)	5.036*** (1.037)	4.861*** (1.034)
Impatience		-2.910* (1.514)	-3.677*** (1.062)	-3.774*** (1.045)	-4.556*** (1.059)
Prefers work			-1.614 (1.519)		
Works hard				0.395 (1.512)	
Thinks of future					-3.003** (1.512)
<i>House</i> × Wealth	2.628 (1.658)				
<i>Cash</i> × Wealth	-4.549 (3.196)				
<i>House</i> × Impatience		-2.923 (2.177)			
<i>Cash</i> × Impatience		6.147 (4.121)			
<i>House</i> × Prefers work			1.089 (2.170)		
<i>Cash</i> × Prefers work			1.523 (4.119)		
<i>House</i> × Works hard				-2.086 (2.174)	
<i>Cash</i> × Works hard				1.427 (4.120)	
<i>House</i> × Thinks of future					-1.852 (2.167)
<i>Cash</i> × Thinks of future					-3.504 (4.098)
Observations	195,405	96,516	96,516	96,516	96,516
R ²	0.549	0.578	0.578	0.578	0.578

Notes: This table reports OLS estimates of equation (2). The outcome variable is 2019 savings after April 8 as a percentage of the *Control* group mean. Column (1) interacts the *House* and *Cash* indicators with wealth (SISBEN); Columns (2)–(5) interact them with impatience (Column 2) and three self-reported measures of financial myopia (Columns 3–5). Each moderator enters as a binary indicator for being above the sample median. All specifications control for baseline savings and the main effect of the interacted measure (and, from Column 2, for wealth). The sample excludes *House+Cash* and accounts with a missing SISBEN score; columns (2)–(5) further restrict to accounts matched to the survey. Robust standard errors are reported in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

liquidity constraints as a confounding factor.²⁸

Predictive power of survey measures Various survey measures meaningfully predict savings behavior. Wealthier workers, as measured by an above-median SISBEN level, save 3.8–5.0 percentage points more than their less wealthy counterparts, and this relationship is highly significant

²⁸This is important because myopic workers may also be those who are most liquidity constrained, which does not allow them to save or respond to our treatments. Controlling for our wealth proxy separates the pure role of myopia from liquidity constraints. We do not additionally control for variables used to match administrative records to the survey (such as region, age, or gender), because these variables are already embedded in the construction of the imputed survey measures. Conditioning on them in the regression mechanically absorbs part of the variation used to identify heterogeneity and exacerbates attenuation from measurement error. In unreported specifications, additionally conditioning on these matching variables pulls some coefficients further toward zero, as expected, and does not affect our conclusions.


across all specifications ($p = 0.001$). Impatience is negatively associated with savings. Workers above the median discount rate save around 2.9 percentage points less than their more patient peers ($p = 0.055$). Among the personality traits in Table 1, forward-looking behavior exhibits the strongest predictive power: workers who report thinking about how their actions affect their future save 3.0 percentage points less than those who do not ($p = 0.047$). The remaining personality measures do not significantly predict savings.

Targeting Despite this meaningful heterogeneity in baseline savings behavior, we find no evidence that lottery treatments differentially affect any of these subgroups. The interaction terms between treatment assignment and our heterogeneity measures are statistically insignificant, and even directionally these estimated effects do not consistently point to beneficial targeting. These findings have important implications for the targeting properties of savings lotteries. A paternalistic policymaker might hope that lottery incentives would disproportionately benefit present-biased individuals, precisely those whose low savings rate is most attributable to self-control problems rather than rational optimization. Our results provide no evidence of such targeting.

Distributional effects While treatment effects do not vary significantly, lotteries may still benefit workers differentially if qualification rates for the lottery vary with wealth. In Figure 6, we plot the average number of qualifications for a lottery prize across the distribution of the wealth index. In Panel A, we find that the average number of house lottery qualifications earned by workers with a wealth index below 60 is 0.044 out of 2, while it is about 3.3 times as large (0.145 out of 2) for workers with wealth index above 75. Panel B shows a similar pattern for the cash lotteries: Among workers with a wealth index below 60, the average number of qualifications is around 0.287 out of 18, but goes up to around 0.868 out of 18 for index values above 75. Since the prizes are disproportionately paid to accounts with a higher wealth index, this implies that both savings lotteries in our study are regressive. This highlights a crucial lesson for policy design: because these lotteries reward absolute *levels* of savings rather than *changes* in savings, they tend to favor wealthier workers who have higher savings in general. An alternative design that rewards marginal improvements in savings relative to an individual reference point (e.g., past savings) could in principle overcome this regressive feature.

Result 5: *Lotteries disproportionately benefit wealthier workers and do not detectably target present-biased workers.*

Figure 6: Qualification for Lottery Prize by Wealth



Notes: This figure presents the average number of lottery qualifications across the wealth index (the worker’s SISBEN score—Colombia’s means-test index for social-program eligibility, with higher values denoting wealthier workers). Panel A plots the average number of house lotteries an account qualified for in 2019 (out of a possible 2), defined as the sum of qualification for the first and second house lottery. Analogously, Panel B plots the average number of cash lottery qualifications in 2019 (out of a possible 18), defined as the number of biweekly periods in which an account qualified. The figure includes all four experimental arms, pooling full- and partial-dose accounts.

3 Welfare Implications of Savings Lotteries

Our reduced-form findings show that savings lotteries generate limited savings effects and crowd out insurance coverage. Whether these interventions raise or reduce welfare, however, is a quantitative question that depends on how much workers value the expected lottery prize and how much they value life insurance coverage. Answering it requires knowledge of the curvature of the utility function to measure risk aversion. While a lottery experiment could in principle be used to recover risk preferences, it would require that subjects know the odds of winning the lottery. Different from lotteries with given odds, however, our savings lotteries involve odds that are *endogenous* to how many workers have met the threshold (the prize is raffled among everyone who qualifies). A worker may systematically misperceive how many others qualify, in which case we cannot assume that their behavior reflects savings responses under *objectively correct* odds. The challenge is, therefore, to separately identify (i) risk preferences and (ii) workers’ *subjective* odds.

As we discuss further below, we recover the curvature of the utility function using a complementary nationwide experiment from a companion study (Gars et al., 2026) that offered *deterministic* savings matches to workers from the same population. Given the estimated preferences from this study, we can then isolate the subjective odds that rationalize workers’ savings response to the lottery. This is a revealed preference approach that recovers subjective beliefs directly from behavior (Rodemeier, 2025; Miller, Sahni, and Strulov-Shlain, 2026).

For context, we first report the relevant reduced-form results of the matching experiment. We then show how they inform parameters of a simple model that quantifies the welfare effects of the lotteries.

3.1 *Complementary Experiment on Deterministic Matches*

During February and March of 2022, trained employees at the Colpensiones call center contacted account holders listed in BEPS administrative records and explained the savings-match incentive. Accounts were randomly assigned to receive a savings match m of 2,500, 5,000, or 10,000 COP if they saved a positive amount within two months of the call. If treated subjects made a deposit during the qualification period, the matched amount was added to their retirement account. A *Control* group received no match but was also called and given the same general information about the BEPS program. To improve comparability, we restrict our analysis to accounts that took part in the lottery experiment in 2019, so that our inputs to the welfare analyses are obtained from a comparable sample of accounts. As with any cross-study comparison, the caveat of drawing comparisons between the two experiments is that they occurred at different points in time.²⁹

Appendix Table A4 reports the effects on the probability of saving any positive amount, on total savings during the qualification period (in % of the *Control* mean), and on post-treatment savings in 2022. A 2,500 COP match raises the savings probability by 12.4 percentage points ($p < 0.001$), a relative effect of 50.2% over *Control*. Increasing the matching amount to 5,000 and 10,000 has a small incremental effect of 1.3 and 2.3 percentage points, respectively. Total qualification-period savings rise by 32.2–44.9% net of the mechanical match transfers.³⁰ Unlike the lotteries, matches generate no offsetting decline afterward: post-treatment savings remain 20.6–27.6% higher.

To put the savings effects on a common economic scale, we compare instruments by their cost-effectiveness. We define cost effectiveness as the change in savings for each peso spent by Colpensiones on the intervention. A 2,500 COP match raises workers' savings by about 15 COP per peso spent by Colpensiones, against roughly 2 COP for the house lottery and 0.1 COP for the cash lottery (Appendix Figure A2). This suggests that the matching intervention is more cost-effective, but we read this benchmarking cautiously, because the two instruments differ in more than just their stochasticity (e.g., in the position of the qualifying savings threshold). We therefore treat the matching experiment primarily as the source of the structural inputs for the welfare analysis to which we now turn.

²⁹Comparisons across studies are increasingly common in public policy evaluations (see, e.g., Hendren and Sprung-Keyser, 2020).

³⁰Choukhmane et al. (2024) show that in the U.S. context, employer matches raise retirement wealth primarily through mechanical transfers, with comparatively modest behavioral responses. Relative to this benchmark, the behavioral effects we measure are large.

3.2 Framework

The goal is a compact model that takes our reduced-form treatment effects as inputs and delivers a first-order approximation of the welfare effect of the lottery. To keep the model transparent, we evaluate the welfare consequences of a worker's *first* interaction with a lottery rather than the full sequence of eighteen cash and two house draws. We do this because savings responses to later lotteries are likely confounded by the responses to earlier lotteries. Valuing these later draws would require taking a stand on how much of a worker's liquid resources have been drawn down by earlier qualification responses, how beliefs have already been updated by prior draws, and on other state variables. Disciplining these state variables becomes intractable and obscures the mechanisms driving the results. Focusing on the first lottery offer keeps the mapping from the observed response to beliefs and welfare both tractable and interpretable, at the cost of silencing any welfare effects that operate only through repeated exposure.

In period 0, the worker earns income y_0 and chooses illiquid savings s_0 , yielding contemporaneous utility $u_0(y_0 - s_0)$.³¹ Savings earn a gross return of $R = 1 + r$. Because the lotteries in our experiment run on qualification windows shorter than a year, a model period corresponds to a fraction $1/f$ of a year, with f the number of periods per year.

Workers enter retirement in period T . Throughout, we use a single object to denote retirement-period wealth: for any savings path $\{s_0, \dots, s_{T-1}\}$,

$$W_T(s_0, \dots, s_{T-1}) = y_T + \sum_{t=0}^{T-1} s_t R^{T-t},$$

where y_T is any income obtained during retirement, such as from government support programs. In the presence of a savings lottery, period-0 savings increase the probability of winning the lottery in period 1. The worker earns lottery tickets $\tau = s_0/p$, where p is the savings required per ticket. In our experiment $p \in \{55,000; 500,000\}$ with corresponding prizes $x \in \{100,000; 102,000,000\}$. If the account wins, the worker's period-1 utility is $u_1(y_1 - s_1 + x)$; otherwise it is $u_1(y_1 - s_1)$.

For a lottery with J participants, the winning probability of subject i is $\omega(\tau_i) = \tau_i / \sum_j \tau_j$, so each additional ticket increases both the numerator and the denominator. This is a Tullock contest in which each player's optimal strategy depends on the aggregate strategies of all other players (Tullock, 1980). In a large contest like ours, with many participants and small individual ticket holdings, each player is effectively of measure zero (Lahkar and Mukherjee, 2023), so the winning probability is approximately linear in tickets, $\omega(\tau_i) \approx \pi \tau_i$, with π denoting the winning probability per ticket.³² We allow workers to have a *subjective* winning probability $\hat{\pi}$ that may differ from the

³¹ Since our primary source of identification will come from relatively short-term savings responses to the treatments, we do not explicitly model future income shocks and random fluctuations in assets and wealth over the life-cycle.

³² To see this, let $Q_{-i} = \sum_{j \neq i} \tau_j$ denote the total tickets held by all participants other than i . A first-order Taylor

objective π .

To accommodate paternalistic policy motives common to retirement policies, we further allow for quasi-hyperbolic preferences with discount factor δ and present-focus parameter $\beta \in (0, 1]$, where $\beta < 1$ implies that workers discount future utility due to a preference for instant gratification (Laibson, 1997). We make the standard assumption that the worker is naive about her present bias, i.e., she incorrectly believes that her future consumption-savings choices will be time-consistent. The framework also admits the time-consistent case $\beta = 1$.

The worker chooses savings to maximize lifetime utility:

$$\max_{\mathbf{s}} \left\{ u_0(y_0 - s_0) + \beta \left(\delta \cdot \mathbb{E}_{\hat{\omega}(\hat{\pi}, \frac{s_0}{p})} [u_1(y_1, s_1, x)] + \delta^2 V(\mathbf{s}) \right) \right\} \quad (3)$$

where the period-1 expectation $\mathbb{E}_{\hat{\omega}(\hat{\pi}, \frac{s_0}{p})} [u_1(y_1, s_1, x)] = u_1(y_1 - s_1 + x) \hat{\pi} \frac{s_0}{p} + u_1(y_1 - s_1)(1 - \hat{\pi} \frac{s_0}{p})$ is taken over the *subjective* probabilities, $\hat{\omega}(\hat{\pi}, \frac{s_0}{p}) = \hat{\pi} \frac{s_0}{p}$. The continuation value is $V = \sum_{t=2}^{T-1} \delta^{t-2} u(y_t - s_t) + \delta^{T-2} u_T(y_T + \sum_{t=0}^{T-1} s_t R^{T-t})$. Let s_0 denote period-0 savings without the lottery and s_0^l savings with the lottery, with corresponding retirement wealth $W_T^0 \equiv W_T(s_0, \dots)$ and $W_T^l \equiv W_T(s_0^l, \dots)$. The first-order condition in the presence of the lottery is

$$\underbrace{-u'_0(y_0 - s_0^l)}_{\text{cost of consumption shifting}} + \underbrace{\beta \delta \cdot \frac{\hat{\pi}}{p} [u_1(y_1 - s_1^l + x) - u_1(y_1 - s_1^l)]}_{\text{perceived lottery gain}} + \underbrace{\beta \delta^T R^T u'_T(W_T^l)}_{\text{savings benefit at retirement}} = 0, \quad (4)$$

and in the absence of the lottery

$$-u'_0(y_0 - s_0) + \beta (\delta R)^T u'_T(W_T^0) = 0. \quad (5)$$

The experiment shifts the worker's FOC from equation 5 to equation 4 by making her aware of the lottery. The implied treatment effect $\Delta s = s_0^l - s_0$ is what we observe empirically. As in other structural exercises (e.g., Abaluck and Gruber, 2011; Cohen and Einav, 2007), we parameterize utility as CARA, $u(c) = -e^{-\gamma c}/\gamma$ with γ measuring absolute risk aversion.

Normative welfare The social planner evaluates normative utility, which differs from the worker's perceived utility for several reasons. We make the common assumption that the social planner treats time-inconsistent behavior as a mistake, so the planner evaluates welfare with $\beta = 1$. The planner also uses the objectively correct winning probability π rather than the potentially biased $\hat{\pi}$ when evaluating the value of the lottery. We focus on a representative agent to keep the analysis transpar-

expansion of $\omega(\tau_i) = \tau_i/(\tau_i + Q_{-i})$ around $\tau_i = 0$ gives $\omega(\tau_i) = \tau_i/Q_{-i} + O(\tau_i^2/Q_{-i}^2)$. In our setting, the national ticket pool is approximately 7,400 for the first cash lottery and 10,400 for the first house lottery, while individual holdings rarely exceed one ticket, so the approximation error is of order 10^{-4} or smaller.

ent. Allowing for heterogeneity would be important if lotteries selectively targeted present-biased workers, which is not what we find in our survey evidence (Section 2.5). We further include in the social welfare function the loss in insurance coverage induced by the lottery.

Since we find no evidence of long-run effects of any of the lottery interventions, they did not change the share of workers who will eventually qualify for the non-contributory retirement safety net, *Colombia Mayor* (recall Section 1.1). This means the lotteries do not affect the fiscal externality associated with undersaving, so we abstract from this channel and focus on the relevant behavioral frictions in our welfare computations.

The social welfare function can then be written as

$$W = u_0(y_0 - s_0(\beta, \hat{\pi})) + \delta \cdot \left[\pi \frac{s_0(\beta, \hat{\pi})}{p} u_1(y_1 - s_1 + x) + \left(1 - \pi \frac{s_0(\beta, \hat{\pi})}{p}\right) u_1(y_1 - s_1) \right] + \sum_{t=2}^{T-1} \delta^t u(y_t - s_t(\beta, \hat{\pi})) + \delta^T u_T(W_T) + \mathcal{L} \quad (6)$$

where $\mathcal{L}$ is the discounted value of life insurance coverage offered by BEPS.

The experimental intervention shifts beliefs from $\hat{\pi}_0 = 0$ to some $\hat{\pi}_1 > 0$, inducing treatment effects Δs_t . Empirically, we saw that the first lottery weakly increased savings only during the qualification period ($\Delta s_0 \geq 0$) and weakly decreased savings only immediately afterwards ($\Delta s_1 \leq 0$). Consistent with this, we set $\Delta s_t = 0$ for $t = 2, \dots, T-1$. Because the framework evaluates a *single* lottery offer, this restriction only requires that the response to an offer is confined to its qualification window and the period immediately after, rather than propagating into long-run future savings. This assumption places no restriction on responses to subsequent lottery offers, which lie outside our model.³³ Therefore, the only welfare-relevant periods are 0, 1, and T. As we prove in Appendix B.3, a first-order approximation of the welfare effect of the lottery is then obtained by:³⁴

$$\Delta W \approx \underbrace{\rho}_{\text{expected prize}} + \Delta s_0 \left[\underbrace{\frac{\delta \Delta u_1}{p} (\pi - \beta \hat{\pi})}_{\text{belief distortion}} + \underbrace{\delta^T u'_T R^T (1 - \beta)}_{\text{undersaving correction}} \right] + \underbrace{\Delta s_1 \delta^T u'_T R^{T-1} (1 - \beta)}_{\text{intertemporal spillover}} - \underbrace{\Delta \mathcal{L}}_{\text{insurance loss}} \quad (7)$$

³³This is supported by the absence of habit formation documented in Section 2: even for the house lottery, where the contemporaneous response is sizable, treatment effects show no persistence. The first cash lottery had zero effect during the qualification period and, therefore, cannot induce any persistence. Savings responses in the Cash treatment later in the year reflect contemporaneous reactions to *later* cash lotteries rather than dynamic effects of the first offer.

³⁴A first-order approximation avoids the additional assumptions and data requirements of a fully specified life-cycle model, which would require information on long-run asset holdings outside the BEPS system that are not available in our data. Factors such as income shocks and liquidity constraints that would be important in a life-cycle model have limited bearing on this analysis, which relies on savings in periods 0 and 1 and consumption only again at retirement. Due to its simplicity and transparency, welfare computations that rely on first-order approximations of this kind are standard in empirical public finance.

where $\rho = \pi \cdot (s_0/p) \cdot (\delta/\gamma)[e^{-\gamma c_1} - e^{-\gamma(c_1+x)}]$ is the expected prize gain and $\Delta u_1 = [e^{-\gamma(y_1-s_1^l)} - e^{-\gamma(y_1-s_1^l+x)}]/\gamma$. We let $\Delta\mathcal{L}$ denote the welfare loss from reduced insurance qualification, due to the lottery-induced reduction in the qualification probability.

If workers are time-consistent and subjective odds are accurate ($\beta = 1$ and $\hat{\pi} = \pi$), the belief-distortion, undersaving-correction, and intertemporal spillover terms vanish by the envelope theorem.³⁵

3.3 Calibration and Identification

Evaluating equation (7) requires estimates of (i) time preferences (β, δ), (ii) risk aversion (γ), (iii) subjective winning probabilities ($\hat{\pi}$), (iv) the welfare value of insurance loss ($\Delta\mathcal{L}$), and (v) the empirical treatment effects ($\Delta s_0, \Delta s_1$). We discuss each in turn.

To keep the exposition concise, we focus the welfare analysis on the *House* lottery, which has a longer response window and yields more precisely estimated savings effects than the *Cash* lottery. The *Cash* lottery yields the same policy conclusion and is deferred to the Appendix.

Time preferences Following Laibson et al. (2026), we calibrate $\beta = 0.5$ and the discount factor to compound to an *annual* value of $\delta^f = 0.99$. Gross annual returns of the retirement account are calibrated to its historical performance of 5%, i.e., $R^f = 1.05$.

Estimation of curvature of utility function We estimate the curvature of the utility function, γ , using data from the complementary matching experiment for the same workers that were part of the lottery experiment. The matching experiment incentivized subjects to increase their probability of saving any positive amount during the qualification period. Let m denote the size of the match and $s_0(m)$ the chosen savings during the qualification period. As we prove in Appendix B.2, observing how the probability of saving changes in response to different match levels identifies γ . Specifically,

$$\gamma \approx \frac{-\frac{\partial \ln\left(\frac{\partial \Pr(s_0(m)>0)}{\partial m}\right)}{\partial m}}{(W'_T(m) + 1)} \quad (8)$$

The numerator measures how the increase in the savings probability changes as we raise the match level. Intuitively, the curvature of this “extensive margin savings function” is directly pinned down by the curvature of the utility function. The denominator $\frac{\partial W_T(m)}{\partial m}$ is the change in retirement wealth due to the increased savings response to the match, which we observe from the treatment

³⁵In principle, a paternalistic planner could eliminate distortions through “optimal deception” by designing information so that $\hat{\pi} = [\pi + \delta^T u'_T R^T (1 - \beta)p / (\delta \Delta u_1)] / \beta$. We do not explore this possibility here.

effects.³⁶ We estimate a value of $\gamma = 0.0008578$ per 1,000 COP ($p < 0.001$) in Appendix B.2, which translates to around 0.002814 per USD and falls into the range of estimates in the literature.³⁷

Subjective winning probabilities Note that our two experiments give us savings responses to a stochastic and to a deterministic incentive. If workers make consistent choices, we can use the savings response to the deterministic match to recover the subjective odds $\hat{\pi}$ that rationalize the observed savings response to the lottery. This identification strategy is analogous to the static setting in Rodemeier (2025) but applied to a dynamic savings-consumption context. As before, subjective probabilities are estimated from the savings response to the *first* lottery within each treatment—which isolates the incentive effect of the lottery from other compounding mechanisms over time. The full identification argument, the calibration of secondary parameters, and the estimation procedure are described in Appendix C. The key estimates are as follows. For the first cash lottery, $\hat{\pi} \approx 0.0\%$ while the true probability is $\pi = 2.7\%$. For the first house lottery, $\hat{\pi} \approx 1.3\%$ against a true probability $\pi = 0.0096\%$, an overestimation by about two orders of magnitude. This indicates that workers underreact to low-prize lotteries and overreact to high-prize lotteries, consistent with probability weighting (Kahneman and Tversky, 1979) and more recent theories where promotion features—like probabilities and prizes—compete for attention (Bordalo et al., 2012; Gennaioli et al., 2026; Kőszegi and Szeidl, 2013). This pattern is robust across a wide grid of alternative parameter calibrations (see Appendix C.1).

Value of insurance coverage Under CARA utility, a worker with annual income y facing mortality risk q and death benefit B has willingness to pay P^* for insurance solving $(1 - q)u(y - P^*) + qu(B) = (1 - q)u(y) + qu(0)$.³⁸ Using our estimated γ , a mortality probability of 0.2% in this population of workers,³⁹ and the BEPS death benefit, we obtain $P^* \approx 1,000,000$ COP (about 305 USD, or 12% of annual income). The expected money-metric welfare loss from reduced coverage is then $\Delta\mathcal{L} = \sum_{n=1}^2 \delta^{fn} \Delta\theta_n P^*$, where $\Delta\theta_n \geq 0$ is the lottery-induced reduction in the coverage

³⁶Specifically, we sum the contemporaneous effect on deposits during the qualification period and the persistent post-treatment effect on deposits documented in Section 3.1. Both components are projected forward to retirement at the gross return R .

³⁷A common calibration for CRRA in the macro and finance literature is a value around 2 (see, e.g., Laibson et al., 2026). The relationship between CARA and CRRA is given by $\sigma = \gamma \cdot W$, where σ is the coefficient of relative risk aversion and W is wealth. Since we do not observe workers' total wealth, we cannot directly translate our CARA estimate into a CRRA coefficient. However, to obtain $\sigma = 2$, workers' accumulated wealth would have to be approximately three-to-four months of average informal wages. While this is not an unreasonable wealth level for many of these economically disadvantaged workers, the total value of assets they hold is likely somewhat higher. This results in an implied CRRA coefficient above 2, as we would expect for poorer workers.

³⁸We assume the family of the deceased has the same utility function and is valued equally by the worker. An alternative specification with an explicit altruism parameter has little impact on our conclusions.

³⁹Derived from Colpensiones insurance-claims data for 2016–2022.

probability in the two respective post-intervention years $n \in \{1, 2\}$.⁴⁰

Choice of lottery and treatment-effect windows For the first house lottery, we set a model period of three months, which is the length of the qualification window in which subjects knew about the lottery (April to end of June). The treatment effect during this period measures $\Delta s_0 > 0$.⁴¹ Consistent with this, we measure $\Delta s_1 < 0$ over three months as well, from July through September. This captures the significant negative spillover effect on savings immediately after the first lottery (recall Figure 3). Our estimate of Δs_1 will weakly *understate* the full negative spillover effect because some subjects in the data may already increase savings in response to the second house lottery, which biases our estimate of Δs_1 upward.⁴² Our conclusion that lotteries reduce welfare would only be strengthened in this case. The cash lottery is treated analogously.⁴³

3.4 Welfare Estimates

Figure 7 shows the direct welfare effects of the house lottery in COP, as well as the indirect effect due to the change in life insurance coverage. The direct effects are: (i) the expected transfer from the lottery prize, (ii) the belief distortion, (iii) the correction in undersaving due to present bias, and (iv) the intertemporal spillover. The figure reports each of these components along with their net total. We find that even though the estimated belief bias is substantive, this bias carries little weight for welfare. Most of the direct welfare effects are due to savings increases and intertemporal spillovers, which partly offset one another, but nonetheless sum up to a (small) positive direct effect of the policy. However, once we take into account the welfare loss from reduced life insurance coverage (−19,166 COP), the welfare implications flip, with a total welfare effect of −14,663 COP.⁴⁴ If we

⁴⁰Coverage in a calendar year is determined by qualification in the prior year (Section 1), so $n = 1$ and $n = 2$ correspond to coverage years 2020 and 2021, driven by the qualification effects in 2019 and 2020 shown in Figure 5. In percentage points, $\Delta\theta_1 = 0.62$ and $\Delta\theta_2 = 1.2$ for the cash lottery, and $\Delta\theta_1 = 0.35$ and $\Delta\theta_2 = 0.85$ for the house lottery. We interpret these reductions as workers forgetting rather than rationally forgoing coverage, given the large willingness to pay estimate we obtain, together with the survey evidence that insurance is a primary motivation to deposit, and—most importantly—the fact that qualifying for the insurance and the lottery was not mutually exclusive.


⁴¹Since savings prior to April also count for the qualification of the house lottery, we compute the number of lottery tickets collected, s_0/p , using all savings between January and June. Thus, we correctly identify the savings response Δs_0 using the time period in which subjects knew about the lottery, while computing the number of earned lottery tickets by using the full qualification period of the first semester of the year.

⁴²We view the importance of this measurement bias as relatively minor since the next house lottery deadline is still several months away, and the reduced-form evidence indicates that workers respond primarily near deadlines.

⁴³For the cash lottery, a model period equals the two-week qualification window, so Δs_1 is measured over the biweek following Δs_0 , which coincides with the next cash lottery’s qualification period. As with the house lottery, this overlap can only *understate* the negative spillover and thus strengthens our welfare conclusion. Here it is also quantitatively irrelevant, since the cash treatment effects over these first periods are statistically indistinguishable from zero throughout.

⁴⁴Since we are only looking at welfare effects on BEPS account holders, we are ignoring how the lottery prize is financed through revenue from taxpayers outside the BEPS system. Including financing costs would further reduce the welfare benefits.

Figure 7: Welfare Effects of the House Lottery



Notes: This figure presents the welfare effects of the house lottery in COP, decomposed into direct and indirect components. The direct components are: (1) the expected prize transfer (ρ), (2) the welfare effect of belief distortions arising from misperceptions of the probability of winning, (3) the welfare gain from correcting undersaving driven by present bias, and (4) the welfare loss from intertemporal spillover of savings across periods. The sum of these four components is reported as the Direct net. The indirect component captures the reduction in life and disability insurance coverage ($\Delta\mathcal{L}$), calculated as the present discounted value of foregone insurance benefits. The sum of the Direct net and the indirect component is reported as the Total net.

estimate this model for the cash lottery, the overall negative effect that we find is even larger and almost entirely due to the insurance channel (see Appendix Figure A3).

It is also important to note that these are the ex-ante welfare effects on the *average* worker. Since the lotteries reduced insurance qualification by at most 1.2 percentage points, looking at average welfare effects masks how much the small share of workers who actually lose insurance will suffer. To put this in perspective, the annual value of life insurance (willingness to pay) is around 1,000,000 COP, approximately 12% of annual income. This highlights that the lottery creates a small risk of an enormous welfare loss for individual workers.

Ultimately, this quantification offers the methodological takeaway that a complete policy analysis must track behavior across both time and adjacent domains. Ignoring these effects in our setting would not only overstate the policy's net welfare effect but even reverse policy implications.

4 Conclusion

This paper provides large-scale experimental evidence on Pigouvian lotteries as a tool to encourage retirement savings. Partnering with Colombia’s public pension fund, we show that lotteries substantially increase deposits during qualification periods but that these increases are offset by subsequent reductions, leaving close to no net effect on savings over the medium or long run. We also document important unintended consequences: workers reduce deposit frequency, which results in losing access to valuable life and disability insurance, and the incentives disproportionately benefit the relatively wealthy. Lotteries also do not detectably target present-biased workers who would benefit the most from saving more. Accounting for these channels, we find that lotteries reduce welfare.

While these results uncover the pitfalls of existing Pigouvian lotteries, it would be too soon to discard lotteries from the public finance toolbox. Our results point to several promising directions. Future work should explore alternative designs of lotteries, such as larger but less frequent prizes, payment structures that reward changes rather than levels of savings, closer integration with existing savings incentives, and more direct targeting of present-biased individuals.

More broadly, our findings highlight that behavioral interventions cannot be evaluated on the targeted margin alone. Localized and short-run effects can mask offsetting responses over time, spillovers onto complementary programs, and regressive incidence. Tracking households across domains and years, together with benchmarking against traditional policy tools, is necessary to determine whether behavioral public policies improve welfare.

References

- Abaluck, J. and J. Gruber (2011). Choice inconsistencies among the elderly: evidence from plan choice in the medicare part d program. *American Economic Review* 101(4), 1180–1210.
- Agence France-Presse (2021). Hong Kong tycoons offer HK\$10.8 million flat to boost Covid job take-up. Hong Kong Free Press.
- Altmann, S., A. Grunewald, and J. Radbruch (2022). Interventions and cognitive spillovers. *Review of Economic Studies* 89(5), 2293–2328.
- Atalay, K., F. Bakhtiar, S. Cheung, and R. Slonim (2014). Savings and prize-linked savings accounts. *Journal of Economic Behavior & Organization* 107, 86–106.
- Barberis, N. (2013). The psychology of tail events: progress and challenges. *American Economic Review* 103(3), 611–616.
- Bauer, R., I. Eberhardt, and P. Smeets (2022). A fistful of dollars: Financial incentives, peer information, and retirement savings. *Review of Financial Studies* 35(6), 2981–3020.
- Becker, G. S. and K. M. Murphy (1988). A theory of rational addiction. *Journal of Political Economy* 96(4), 675–700.
- Beshears, J., J. J. Choi, D. Laibson, and B. C. Madrian (2018). Behavioral household finance. In *Handbook of behavioral economics: Applications and foundations 1*, Volume 1, pp. 177–276. Elsevier.
- Beuermann, D. W., B. Hoffmann, M. Stampini, D. L. Vargas, and D. Vera-Cossio (2025). Shooting a moving target: Evaluating targeting tools for social programs when income fluctuates. *Journal of Development Economics* 172, 103395.
- Bordalo, P., N. Gennaioli, and A. Shleifer (2012). Salience theory of choice under risk. *Quarterly Journal of Economics* 127(3), 1243–1285.
- Browne, K. (2016). California cities shine in \$5 million energy prize challenge. Western City Magazine.
- Burke, J. (2021). Do prize-linked incentives promote positive financial behavior? evidence from a debt reduction intervention. *Journal of Public Economics* 204, 104534.
- Campos-Mercade, P., A. N. Meier, S. Meier, D. G. Pope, F. H. Schneider, and E. Wengström (2024). Incentives to vaccinate. Technical report, National Bureau of Economic Research.

- Charness, G. and U. Gneezy (2009). Incentives to exercise. *Econometrica* 77(3), 909–931.
- Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics* 1(1), 451–488.
- 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. *Quarterly Journal of Economics* 129(3), 1141–1219.
- Choukhmane, T. (2025). Default options and retirement saving dynamics. *American Economic Review* 115(11), 3749–3787.
- Choukhmane, T., J. Colmenares, C. O’Dea, J. L. Rothbaum, and L. D. Schmidt (2024). Who benefits from retirement saving incentives in the US? evidence on gaps in retirement wealth accumulation by race and parental income. Technical report, National Bureau of Economic Research.
- Choukhmane, T. and C. Palmer (2024). How do consumers finance increased retirement savings? Unpublished working paper.
- Cohen, A. and L. Einav (2007). Estimating risk preferences from deductible choice. *American economic review* 97(3), 745–788.
- Cole, S., B. Iverson, and P. Tufano (2022). Can gambling increase savings? empirical evidence on prize-linked savings accounts. *Management Science* 68(5), 3282–3308.
- Colpensiones (2023). Efectos de la entrega de los incentivos de bono de vivienda y motivaciones para ahorrar y participar en los sorteos BEPS. Technical report.
- Cookson, J. A. (2018). When saving is gambling. *Journal of Financial Economics* 129(1), 24–45.
- DellaVigna, S. (2018). Structural behavioral economics. In *Handbook of behavioral economics: Applications and foundations 1*, Volume 1, pp. 613–723. Elsevier.
- Dupas, P. (2014). Short-run subsidies and long-run adoption of new health products: Evidence from a field experiment. *Econometrica* 82(1), 197–228.
- Falk, A., A. Becker, T. Dohmen, D. Huffman, and U. Sunde (2022). The preference survey module: A validated instrument for measuring risk, time, and social preferences. *Management Science* 69(4), 1935–1950.
- Filiz-Ozbay, E., J. Guryan, K. Hyndman, M. Kearney, and E. Y. Ozbay (2015). Do lottery payments induce savings behavior? evidence from the lab. *Journal of Public Economics* 126, 1–24.

- Gars, J., L. Prada, E. Tripodi, and S. Borda (2025). Personalized reminders: Evidence from a field experiment with voluntary retirement savings in Colombia. *Journal of Behavioral and Experimental Finance* 48, 101112.
- Gars, J., M. Rodemeier, and E. Tripodi (2026). Saving attention. mimeo.
- Gennaioli, N., M. Rodemeier, and G. Sun (2026). Nonlinear pricing and revealed salience.
- Gertler, P., S. Higgins, A. Scott, and E. Seira (2023). Using lotteries to attract deposits. Technical report, National Bureau of Economic Research.
- Gneezy, U., S. Meier, and P. Rey-Biel (2011). When and why incentives (don't) work to modify behavior. *Journal of Economic Perspectives* 25(4), 191–210.
- Gomes, F., M. Haliassos, and T. Ramadorai (2021). Household finance. *Journal of Economic Literature* 59(3), 919–1000.
- Guillén, M. F. and A. E. Tschoegl (2002). Banking on gambling: Banks and lottery-linked deposit accounts. *Journal of Financial Services Research* 21, 219–231.
- Guryan, J. and M. S. Kearney (2010, August). Is lottery gambling addictive? *American Economic Journal: Economic Policy* 2(3), 90–110.
- Hanna, R., S. Mullainathan, and J. Schwartzstein (2014). Learning through noticing: Theory and evidence from a field experiment. *Quarterly Journal of Economics* 129(3), 1311–1353.
- Hendren, N. and B. Sprung-Keyser (2020). A unified welfare analysis of government policies. *The Quarterly journal of economics* 135(3), 1209–1318.
- Hussam, R., A. Rabbani, G. Reggiani, and N. Rigol (2022). Rational habit formation: experimental evidence from handwashing in India. *American Economic Journal: Applied Economics* 14(1), 1–41.
- Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica* 47(2), 263–291.
- Kearney, M. S., P. Tufano, J. Guryan, and E. Hurst (2010). Making savers winners: An overview of prize-linked savings products. Technical Report 16433, National Bureau of Economic Research.
- Kőszegi, B. and A. Szeidl (2013). A model of focusing in economic choice. *Quarterly Journal of Economics* 128(1), 53–104.

- Kuhn, P., P. Kooreman, A. Soetevent, and A. Kapteyn (2011). The effects of lottery prizes on winners and their neighbors: Evidence from the dutch postcode lottery. *American Economic Review* 101(5), 2226–2247.
- Lahkar, R. and S. Mukherjee (2023). Optimal large population tullock contests. *Oxford Open Economics* 2, odad003.
- Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics* 112(2), 443–477.
- Laibson, D., S. C. Lee, P. Maxted, A. Repetto, and J. Tobacman (2026). Estimating discount functions with consumption choices over the lifecycle. *Review of Financial Studies* 39(6), 1580–1610.
- Landry, C. E., A. Lange, J. A. List, M. K. Price, and N. G. Rupp (2006). Toward an understanding of the economics of charity: Evidence from a field experiment. *Quarterly Journal of Economics* 121(2), 747–782.
- Lange, A., J. A. List, and M. K. Price (2007a). A fundraising mechanism inspired by historical tontines: Theory and experimental evidence. *Journal of Public Economics* 91(9), 1750–1782.
- Lange, A., J. A. List, and M. K. Price (2007b). Using lotteries to finance public goods: Theory and experimental evidence. *International Economic Review* 48(3), 901–927.
- Lockwood, B. B., H. Allcott, D. Taubinsky, and A. Sial (2024, 08). What drives demand for state-run lotteries? evidence and welfare implications. *Review of Economic Studies* 92(4), 2578–2623.
- Medina, P. C. (2021). Side effects of nudging: Evidence from a randomized intervention in the credit card market. *Review of Financial Studies* 34(5), 2580–2607.
- Medina, P. C. and M. Pagel (2025). Does saving cause borrowing? implications for the coholding puzzle. *Journal of Finance* 80(5), 2689–2738.
- Miller, K. M., N. S. Sahni, and A. Strulov-Shlain (2026). Sophisticated consumers with inertia: Long-term implications from a large-scale field experiment. *American Economic Review*. Forthcoming.
- Moleskis, M. (2023). Incentivizing off-peak public transit with traffic lottery. Green Nudges (blog).
- Naritomi, J. (2019). Consumers as tax auditors. *American Economic Review* 109(9), 3031–3072.
- Norsk Tipping (2021). Pantelotteriet. Website. Norwegian bottle deposit lottery run in partnership with the Red Cross.


- OECD (2024). *OECD Reviews of Labour Market and Social Policies: Colombia 2024*. OECD Publishing.
- OECD (2025). *Pensions at a Glance 2025: OECD and G20 Indicators*. Paris: OECD Publishing.
- Pigou, A. C. (1920). *The Economics of Welfare*. London: Macmillan and Co.
- Rodemeier, M. (2025). Buy baits and consumer sophistication: Field evidence from instant rebates. *American Economic Journal: Economic Policy* 17(2), 30–59.
- Rodemeier, M. and A. Löschel (2025). Information nudges, subsidies, and crowding out of attention: Field evidence from energy efficiency investments. *Journal of the European Economic Association* 23(4), 1397–1430.
- Royer, H., M. Stehr, and J. Sydnor (2015). Incentives, commitments, and habit formation in exercise: evidence from a field experiment with workers at a fortune-500 company. *American Economic Journal: Applied Economics* 7(3), 51–84.
- 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.
- Schultz, J. (2010). Speed camera lottery wins VW fun theory contest. The New York Times.
- Somville, V. and L. Vandewalle (2018). Saving by default: Evidence from a field experiment in rural India. *American Economic Journal: Applied Economics* 10(3), 39–66.
- Team Impact (2025). Indian railways turns tickets into lottery entries to curb fare evasion. Impact Online.
- Toubia, O., E. Johnson, T. Evgeniou, and P. Delquié (2013). Dynamic experiments for estimating preferences: An adaptive method of eliciting time and risk parameters. *Management Science* 59(3), 613–640.
- Tufano, P. (2008). Saving whilst gambling: An empirical analysis of UK premium bonds. *American Economic Review* 98(2), 321–326.
- Tullock, G. (1980). Efficient rent seeking. In J. M. Buchanan, R. D. Tollison, and G. Tullock (Eds.), *Toward a Theory of the Rent-Seeking Society*, pp. 97–112. Texas A&M University Press.
- United States Congress (2014). American savings promotion act. Public Law No: 113-251.
- Vazquez-Bare, G. (2023). Identification and estimation of spillover effects in randomized experiments. *Journal of Econometrics* 237(1), 105237.

- Villa, J. M. and A. Benitez (2023). Beneficios económicos periódicos (BEPS): promoviendo el ahorro en un entorno de informalidad. In Tuesta and Bhardwaj (Eds.), *Pensiones para todos: Propuestas para Sistemas de Pensiones Más Inclusivos en Latinoamérica*, pp. 263–294. India PinBox Solutions Pte. Ltd.
- Walkey, A. J., A. Law, and N. A. Bosch (2021). Lottery-based incentive in Ohio and COVID-19 vaccination rates. *JAMA* 326(8), 766–767.
- World Bank (2014). Step skills measurement household survey 2012 (wave 1).

Online Appendix


A Additional Figures and Tables

Figure A1: Bunching at Cash Lottery Threshold over Time




Notes: This figure presents the effects of the *House*, *House+Cash*, and *Cash* treatments on the likelihood of choosing savings equal to the cash lottery qualification threshold for each biweek, based on equation (1). Panels plot the *House* treatment (Panel A), the *House+Cash* treatment (Panel B), and the *Cash* treatment (Panel C). Thick lines denote the intervention periods for the house lottery (Panel A), from April 11 to December 31, 2019, the *House+Cash* lotteries (Panel B), from April 11 to December 31, 2019, and the cash lottery (Panel C), from April 11 to December 15, 2019, respectively. Each coefficient is from a separate regression where the outcome variable indicates if total savings in biweek t is equal to 55,000 COP, controlling for baseline savings, age, gender, and account age, with department fixed effects. Each point is plotted at the midpoint of the two-week window it summarizes; the first such window opens on April 8, three days before messaging begins on April 11. Through June 30, 2019, Panels B and C retain the partial-dose *House+Cash* and *Cash* accounts (messed only April–June 2019); they are excluded from later periods. Shaded bands around the point estimates denote 95% confidence intervals; standard errors are robust.

Figure A2: Cost Effectiveness: Lotteries versus Matches



Notes: This figure presents estimates of the total change in savings per COP spent by Colpensiones for the house and cash lotteries, and three match levels (2,500, 5,000, and 10,000 COP)—a benefit–cost ratio, where a value of X means X COP of additional saving per COP of program spending. Changes in savings for each treatment are estimated against the control group of the respective intervention. Program spending is the fund’s expected outlay per worker: for the lotteries, the expected prize a worker wins across all draws—their tickets times the per-ticket probability of winning times the prize—which is the same expected transfer that enters our welfare analysis; for the matches, the match amount times the share of workers who save. The dashed horizontal line marks break-even (at 1), where the additional savings induced equal what the fund spends. Error bars denote 95% confidence intervals; standard errors are calculated using the delta method.

Figure A3: Welfare Effects: Cash Lottery



Notes: This figure presents the welfare effects of the cash lottery in COP, decomposed into direct and indirect components. The direct components are: (1) the expected prize transfer (ρ), (2) the welfare effect of belief distortions arising from misperceptions of the probability of winning, (3) the welfare gain from correcting undersaving driven by present bias, and (4) the welfare loss from intertemporal spillovers of savings across periods. The sum of these four components is reported as the Direct net. The indirect component captures the reduction in life and disability insurance coverage ($\Delta\mathcal{L}$), calculated as the present discounted value of foregone insurance benefits. The sum of the Direct net and the indirect component is reported as the Total net.

Table A1: Baseline balance across treatment arms

	(1)	(2)	(3)	(4)	(5)	(6)
	All	Control	House	House+Cash	Cash	<i>p</i> -value
Panel A: Baseline savings (2018)						
Savings (1000s COP)	97.594 (0.315)	97.863 (0.632)	97.554 (0.629)	97.567 (0.629)	97.392 (0.628)	0.962
Positive savings (%)	68.020 (0.075)	68.049 (0.150)	67.983 (0.150)	68.022 (0.150)	68.026 (0.150)	0.992
Number of deposits	3.895 (0.009)	3.897 (0.018)	3.878 (0.017)	3.906 (0.017)	3.899 (0.018)	0.700
≥ six deposits (%)	35.935 (0.077)	36.047 (0.154)	35.878 (0.154)	35.929 (0.154)	35.886 (0.154)	0.857
Qualify for insurance (%)	40.245 (0.079)	40.372 (0.158)	40.211 (0.158)	40.241 (0.158)	40.154 (0.158)	0.794
Panel B: Characteristics						
Female (%)	69.756 (0.074)	69.755 (0.148)	69.750 (0.148)	69.760 (0.148)	69.759 (0.148)	1.000
Average age (years)	50.491 (0.017)	50.493 (0.033)	50.496 (0.033)	50.490 (0.033)	50.485 (0.033)	0.996
SISBEN	31.540 (0.028)	31.476 (0.057)	31.532 (0.057)	31.617 (0.057)	31.536 (0.057)	0.376
Age of account (years)	1.874 (0.002)	1.874 (0.003)	1.873 (0.003)	1.875 (0.003)	1.875 (0.003)	0.983
Active	0.500 (0.001)	0.499 (0.002)	0.500 (0.002)	0.501 (0.002)	0.501 (0.002)	0.767
Observations	387,449	96,862	96,862	96,863	96,862	

Notes: This table reports means and standard errors in parentheses for baseline characteristics across treatment arms. Randomization was conducted at the account level. The *Cash* and *House+Cash* arms pool full- and partial-dose accounts. The *p*-value column reports the result of an *F*-test of equality of means across the four treatment arms. SISBEN is a government welfare index score. *Active* indicates whether the account had at least one deposit in the six months prior to randomization.

Table A2: English Translation of Messages

	(1) Second Thursday	(2) Fourth Thursday
<i>Control</i>	Save with Colpensiones BEPS and protect your old age. Call [REDACTED]	Do you want to build a more secure future? Save in the Colpensiones BEPS program. Tel [REDACTED]
<i>House</i>	Earn with Colpensiones BEPS. Save \$500k by [DATE] and participate in the draw for a home. Call [REDACTED]	Want to win with BEPS-Colpensiones? Accumulate savings of \$500k by [DATE] for the draw for a home. Call [REDACTED]
<i>House+Cash</i>	Earn with Colpensiones BEPS. By saving \$55k every two weeks you can earn \$100k in cash. Accumulate \$500k by [DATE] for housing lottery. Call [REDACTED]	Want to win with BEPS-Colpensiones? Save \$55k every two weeks for \$100k in cash and accumulate \$500k by [DATE] for housing lottery. Call [REDACTED]
<i>Cash</i>	Earn with BEPS-Colpensiones. If you save \$55k every two weeks, you can earn \$100k in cash. Call [REDACTED]	Want to win with BEPS-Colpensiones? Save \$55k every two weeks for the \$100k cash draw. Call [REDACTED]

Notes: This table shows English translations of the text messages that were sent in Spanish to workers.

Table A3: Treatment Effects on Savings

	(1) Levels	(2) Poisson
Panel A: All participants		
<i>House</i>	1.599*** (0.552)	0.020** (0.010)
<i>Cash</i>	1.994* (1.067)	0.023 (0.020)
Observations	207,924	207,924
Control mean	69.4	69.4
Panel B: Active		
<i>House</i>	3.110*** (1.032)	0.020** (0.009)
<i>Cash</i>	4.379** (2.017)	0.029 (0.018)
Observations	103,896	103,896
Control mean	132.0	132.0
Panel C: Inactive		
<i>House</i>	0.009 (0.377)	-0.007 (0.054)
<i>Cash</i>	-0.884 (0.645)	-0.109 (0.108)
Observations	104,028	104,028
Control mean	7.2	7.2

Notes: This table reports treatment effects on 2019 savings after April 8, in thousands of COP. Column (1) reports OLS level effects; column (2) reports average proportional effects (semi-elasticities) from Poisson pseudo-maximum-likelihood estimation. *House* and *Cash* are indicators for the respective treatment arms, with *Control* as the omitted category and *House+Cash* excluded. Active and Inactive refer to whether the account made at least one deposit in the six months prior to randomization. All specifications include department fixed effects and control for baseline savings, age, gender, and account age. Robust standard errors are reported in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A4: Treatment Effects on Savings (Matching Experiment)

	(1) Saving Probability in Qualification Period	(2) Savings (% of <i>Control</i> mean)	(3) Post-Treatment Savings (% of <i>Control</i> mean)
Match 2,500 COP	0.124*** (0.015)	35.906*** (11.987)	22.871** (9.144)
Match 5,000 COP	0.136*** (0.016)	32.177*** (11.555)	20.630** (9.252)
Match 10,000 COP	0.147*** (0.015)	44.906*** (12.755)	27.566*** (9.237)
Observations	7,996	75,762	274,030
Control mean	0.247		

Notes: This table reports treatment effects from the matching experiment. Treatment arms are the deterministic savings matches—2,500, 5,000, or 10,000 COP—with *Control* as the omitted category. Column (1) is an account-level linear probability regression (one observation per account) for saving any positive amount during the qualification period; columns (2)–(3) are account-week panel regressions of savings (relative to the control mean), over the qualification period and the post-qualification period, respectively. All columns control for age, gender, and account age, with department fixed effects; columns (2)–(3) additionally absorb month fixed effects. Robust standard errors in column (1); standard errors clustered by account in (2)–(3). The *Control mean* row reports the control-group saving probability in column (1); the outcomes in columns (2)–(3) are expressed relative to the *Control* mean of savings over the respective period. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B Mathematical Appendix

B.1 Mapping from Subjective Odds to Optimal Savings Response

Proposition 1 (Savings Response to Lottery). *Under CARA preferences, the treatment effect of a lottery with subjective winning probability $\hat{\pi}$ on period-0 savings is approximately:*

$$\Delta s \approx \hat{\pi} \cdot \frac{\frac{\delta}{p} \left[e^{-\gamma(y_1 - s_1^l)} - e^{-\gamma(y_1 - s_1^l + x)} \right] / \gamma}{\gamma \delta^T R^T e^{-\gamma W_T^0} (1 + R^T)}. \quad (9)$$

Proof. The worker maximizes lifetime utility:

$$\max_{\mathbf{s}} \left\{ u_0(y_0 - s_0) + \beta \left(\delta \cdot \mathbb{E}_{\hat{\omega}(\hat{\pi}, \frac{s_0}{p})} [u_1(y_1, s_1, x)] + \delta^2 V(\mathbf{s}) \right) \right\} \quad (10)$$

The First-Order Condition (FOC) for period-0 savings in the presence of the lottery incentive is:

$$-u'_0(y_0 - s_0^l) + \beta \delta \cdot \frac{\hat{\pi}}{p} \underbrace{[u_1(y_1 - s_1^l + x) - u_1(y_1 - s_1^l)]}_{\Delta u_1} + \beta \delta^T R^T u'_T(W_T^l) = 0 \quad (11)$$

In the absence of a lottery ($\hat{\pi} = 0$), the FOC for baseline savings s_0 is:

$$-u'_0(y_0 - s_0) + \beta \delta^T R^T u'_T(W_T^0) = 0 \implies e^{-\gamma(y_0 - s_0)} = \beta \delta^T R^T e^{-\gamma W_T^0} \quad (12)$$

Subtracting equation 12 from equation 11 and using $u'(c) = e^{-\gamma c}$, we obtain:

$$-[e^{-\gamma(y_0 - s_0^l)} - e^{-\gamma(y_0 - s_0)}] + \beta \delta \frac{\hat{\pi}}{p} \Delta u_1 + \beta \delta^T R^T [e^{-\gamma W_T^l} - e^{-\gamma W_T^0}] \approx 0 \quad (13)$$

Applying a first-order Taylor expansion for s_0^l around s_0 and W_T^l around W_T^0 :

$$e^{-\gamma(y_0-s_0^l)} \approx e^{-\gamma(y_0-s_0)} + \gamma e^{-\gamma(y_0-s_0)} \Delta s \quad (14)$$

$$e^{-\gamma W_T^l} \approx e^{-\gamma W_T^0} - \gamma R^T e^{-\gamma W_T^0} \Delta s \quad (15)$$

Substituting these into the difference equation yields:

$$-\gamma e^{-\gamma(y_0-s_0)} \Delta s + \beta \delta \frac{\hat{\pi}}{p} \Delta u_1 - \beta \gamma \delta^T R^{2T} e^{-\gamma W_T^0} \Delta s \approx 0 \quad (16)$$

Next, substitute $e^{-\gamma(y_0-s_0)}$ using the control FOC (equation 12):

$$-\gamma(\beta \delta^T R^T e^{-\gamma W_T^0}) \Delta s + \beta \delta \frac{\hat{\pi}}{p} \Delta u_1 - \beta \gamma \delta^T R^{2T} e^{-\gamma W_T^0} \Delta s \approx 0 \quad (17)$$

which means β cancels out of the equation. Further rearranging yields:

$$\delta \frac{\hat{\pi}}{p} \Delta u_1 \approx \Delta s \left[\gamma \delta^T R^T e^{-\gamma W_T^0} (1 + R^T) \right] \quad (18)$$

Finally, solving for Δs and substituting $\Delta u_1 = [e^{-\gamma(y_1-s_1^l)} - e^{-\gamma(y_1-s_1^l+x)}]/\gamma$ yields the result in equation 9. $\square$

B.2 Identification of Risk Aversion from Matching Experiment

Proposition 2 (CARA Parameter Identification). *Given a matching experiment where workers receive match m for any positive savings, the CARA parameter γ is identified by:*

$$\frac{\partial \ln \left(\frac{\partial \Pr(s_0(m) > 0)}{\partial m} \right)}{\partial m} \approx -\gamma \cdot (W_T^l(m) + 1) \quad (19)$$

Proof. Let V_m denote the value of saving a positive amount with match m , and V_0 the value of not saving. The worker saves if:

$$V_m - V_0 \geq \epsilon \quad (20)$$

where ϵ is an unobserved preference shock with CDF $F(\epsilon)$.

The probability of saving is:

$$S(m) = \Pr(s_0 > 0) = F(V_m - V_0) \quad (21)$$

The marginal value of the match at retirement is:

$$\frac{\partial V_m}{\partial m} = \delta^T u'_T(W_T(m) + m) = \delta^T e^{-\gamma(W_T(m) + m)} \quad (22)$$

Therefore:

$$\frac{\partial S(m)}{\partial m} = \delta^T e^{-\gamma(W_T(m) + m)} f(V_m - V_0) \quad (23)$$

Taking logs:

$$\ln\left(\frac{\partial S(m)}{\partial m}\right) = \ln(\delta^T) - \gamma(W_T(m) + m) + \ln(f(V_m - V_0)) \quad (24)$$

Differentiating with respect to m :

$$\begin{aligned} \frac{\partial \ln\left(\frac{\partial S(m)}{\partial m}\right)}{\partial m} &= -\gamma(W'_T(m) + 1) + \frac{f'(V_m - V_0)}{f(V_m - V_0)} \frac{\partial(V_m - V_0)}{\partial m} \\ &\approx -\gamma(W'_T(m) + 1) \end{aligned} \quad (25)$$

where the approximation requires that $F(\cdot)$ be locally linear—a common requirement in the sufficient statistics literature (Chetty, 2009). $\square$

B.3 Proof of Welfare Effects

Proposition 3 (Welfare Decomposition). *The first-order welfare effect of informing workers about a lottery is:*

$$\begin{aligned} \frac{\partial W}{\partial \hat{\pi}} \Delta \hat{\pi} &= \Delta s_0 \cdot \left(\frac{\delta \Delta u_1}{p} (\pi - \beta \hat{\pi}) + \delta^T u'_T R^T (1 - \beta) \right) \\ &\quad + \sum_{t=1}^{T-1} \Delta s_t \cdot \delta^T u'_T R^{T-t} (1 - \beta) \end{aligned} \quad (26)$$

where $\Delta s_t = \frac{\partial s_t}{\partial \hat{\pi}} \Delta \hat{\pi}$ is the treatment effect on savings in period t .

Proof. Social welfare under the lottery is:

$$\begin{aligned} W(\pi, \mathbf{s}(\beta, \hat{\pi})) &= u_0(y_0 - s_0(\beta, \hat{\pi})) + \delta \cdot \mathbb{E}_{\omega(\pi, s_0)}[u_1(y_1, s_1, x)] \\ &\quad + \sum_{t=2}^{T-1} \delta^t u(y_t - s_t) + \delta^T u_T(W_T) \end{aligned} \quad (27)$$

Taking the total derivative:

$$\frac{dW}{d\hat{\pi}} = \sum_{t=0}^{T-1} \frac{\partial W}{\partial s_t} \frac{\partial s_t}{\partial \hat{\pi}} \quad (28)$$

For $t = 0$:

$$\frac{\partial W}{\partial s_0} = -u'_0(y_0 - s_0) + \delta \cdot \frac{\pi}{p} [u_1(y_1 - s_1 + x) - u_1(y_1 - s_1)] + \delta^T u'_T R^T \quad (29)$$

The worker's FOC implies:

$$u'_0(y_0 - s_0) = \beta \delta \cdot \frac{\hat{\pi}}{p} [u_1(y_1 - s_1 + x) - u_1(y_1 - s_1)] + \beta \delta^T u'_T R^T \quad (30)$$

Substituting:

$$\frac{\partial W}{\partial s_0} = \frac{\delta \Delta u_1}{p} (\pi - \beta \hat{\pi}) + \delta^T u'_T R^T (1 - \beta) \quad (31)$$

For $t \geq 1$:

$$\frac{\partial W}{\partial s_t} = -\delta^t u'(y_t - s_t) + \delta^T u'_T R^{T-t} \quad (32)$$

Using the worker's Euler equation $u'(y_t - s_t) = \beta \delta^{T-t} R^{T-t} u'_T$, we get:

$$\frac{\partial W}{\partial s_t} = \delta^T u'_T R^{T-t} (1 - \beta) \quad (33)$$

Combining all periods yields the result. □

B.4 Welfare Effects of the Lottery Itself

Proposition 4 (Welfare effect of the lottery prize). *The direct welfare effect from the expected prize is:*

$$\Delta W^{prize} = \pi \cdot \frac{s_0}{p} \cdot \frac{\delta}{\gamma} [e^{-\gamma c_1} - e^{-\gamma(c_1+x)}] \quad (34)$$

Proof. A worker saving s_0 obtains s_0/p lottery tickets. Expected utility with the lottery is:

$$\mathbb{E}[U^{\text{lottery}}] = u_0(y_0 - s_0) + \delta [\pi \cdot \frac{s_0}{p} \cdot u_1(c_1 + x) + (1 - \pi \cdot \frac{s_0}{p}) \cdot u_1(c_1)] + V_2 \quad (35)$$

Without the lottery:

$$\mathbb{E}[U^{\text{no lottery}}] = u_0(y_0 - s_0) + \delta u_1(c_1) + V_2 \quad (36)$$

The difference yields:

$$\Delta W^{\text{prize}} = \delta \pi \cdot \frac{s_0}{p} \cdot [u_1(c_1 + x) - u_1(c_1)] \quad (37)$$

Under CARA utility:

$$u_1(c_1 + x) - u_1(c_1) = -\frac{1}{\gamma} e^{-\gamma(c_1+x)} + \frac{1}{\gamma} e^{-\gamma c_1} = \frac{1}{\gamma} [e^{-\gamma c_1} - e^{-\gamma(c_1+x)}] \quad (38)$$

Therefore:

$$\Delta W^{\text{prize}} = \pi \cdot \frac{s_0}{p} \cdot \frac{\delta}{\gamma} [e^{-\gamma c_1} - e^{-\gamma(c_1+x)}] \quad (39)$$

□

C Calibration and Estimation of Subjective Winning Probabilities

In the experiment, we observe the treatment effect $\Delta s = s_0^l - s_0$, which is implicitly pinned down by equating the two FOCs. As proven in Appendix B.1, a first-order Taylor expansion yields:

$$\Delta s \approx \hat{\pi} \cdot \frac{\frac{\delta}{p} [e^{-\gamma(y_1 - s_1^l)} - e^{-\gamma(y_1 - s_1^l + x)}] / \gamma}{\gamma \delta^T R^T u'_T(W_T^0)(1 + R^T)}. \quad (40)$$

This provides a moment condition to estimate $\hat{\pi}$. Notably, β cancels out of this ratio, so the estimates of worker beliefs are invariant to the degree of present focus.

Calibrations. Identification of $\hat{\pi}$ requires knowledge of the treatment effect Δs , along with estimates for income y , baseline savings s_0 , the discount factor δ , retirement wealth W_T^0 , and risk aversion γ . While we do not directly observe y , we observe each worker's social welfare index (SISBEN), which is designed to predict the permanent component of income for determining eligibility for social programs (Beuermann et al., 2025). We transform this distribution into an income distribution of the same shape with a mean equal to the average income of informal workers in 2019. The counterfactual savings without a lottery, s_0 , is not observed for treatment subjects. We therefore use savings during the same time window one year prior to the experiment as a proxy and subtract a general time trend in savings computed from the difference in *Control*-group savings. Because savings are generally small relative to income, measurement error in s_0 has a negligible effect on the estimates. We set the annual discount factor $\delta^f = 0.99$ in line with the literature (Laibson et al., 2026), and verify in Appendix C.1 that our qualitative conclusions are not sensitive to alternative values.

We also require knowledge of marginal utility of consumption at retirement in the absence of

the lottery, $u'_T(W_T^0)$. We make the simplifying assumption that period-0 savings s_0 and its returns are consumed at the beginning of retirement rather than over multiple retirement periods. This simplification captures the fundamental intertemporal trade-off without specifying a multi-period retirement phase. Because s_0 is small, the specific drawdown profile during retirement is not quantitatively important for our estimates. In equation 40, we therefore replace $u'_T(W_T^0)$ by $e^{-\gamma(y_r + s_0(R^T + m))}$, where we project the worker's baseline savings forward over their individual horizon T_i , to their effective claiming age, using a gross real annual return of $R = 1.05$ (a 5% real interest rate, the historical performance of the BEPS portfolio) and the $m = 0.20$ government match on retirement savings. We further add the monthly payment y_r from Colombia's non-contributory retirement safety net *Colombia Mayor* (see Section 1.1), approximately 80,000 COP per month.⁴⁵

Estimation of risk aversion. As we prove in Appendix B.2, observing how the savings probability changes in response to different match levels identifies γ :

$$\frac{\partial \ln \left(\frac{\partial \Pr(s_0(m) > 0)}{\partial m} \right)}{\partial m} \approx -\gamma \cdot (W'_T(m) + 1) \quad (41)$$

To implement this, we approximate the qualification-period saving probability as a quadratic in the match level, $\Pr(s_0 > 0 \mid m) \approx \alpha + a m + b m^2/2$, which gives $\partial \Pr / \partial m = a + b m$ and $\partial \ln(\partial \Pr / \partial m) / \partial m = b / (a + b m)$. Equating this expression with $-\gamma(W'_T(m) + 1)$ at the sample-mean match level identifies γ from the regression coefficients (a, b) and an estimate of the behavioral savings response $\partial s / \partial m$ that enters $W'_T(m)$. We estimate (a, b) by regressing a binary indicator for any deposit during the qualification period on the match level. The behavioral response $\partial s / \partial m$ is obtained from parallel regressions of total qualification-period and post-treatment deposits on the match level; the sum of these two components, projected forward to retirement at the gross return R , captures both the contemporaneous match response and the persistent habit-induced increase in future savings documented in Section 3.1. Delta-method standard errors propagate the sampling uncertainty in the underlying regression coefficients through the nonlinear expression for γ .

Estimation of $\hat{\pi}$. With these primitives at hand, we estimate $\hat{\pi}$ from the data. With treatment assignment $D_i \in \{0, 1\}$, we model savings in period zero as a linear regression, $s_i = \zeta + \hat{\pi} \nu_i D_i + \epsilon_i$,

⁴⁵The horizon T_i runs to each worker's effective claiming age A^* , which we calibrate to 66 for women and 71 for men rather than the statutory minimum retirement ages of 57 and 62. Colombia's average effective labor-market exit ages are 61.2 and 67.4, respectively (OECD, 2025, Figure 6.13), but these are national averages, and informal workers tend to exit later.

with ζ a regression constant and

$$\nu_i = \frac{\frac{\delta}{p} \left[e^{-\gamma(y_i - s_{1i}^l)} - e^{-\gamma(y_i - s_{1i}^l + x)} \right] / \gamma}{\gamma \delta^{T_i} R^{T_i} e^{-\gamma(y_r + s_0(R^{T_i} + m))} (1 + R^{T_i})}$$

Since we observe each subject's age, we compute the individual horizon T_i to the effective claiming age A^* defined above. We estimate $(\zeta, \hat{\pi})$ via GMM by solving the moment condition

$$\mathbb{E}[(s_i - \zeta - \hat{\pi} \nu_i D_i) \cdot (1, D_i)'] = \mathbf{0} \quad (42)$$

With two parameters and two moments, the system is exactly identified. The resulting estimate is interpreted as the subjective per-ticket winning probability that rationalizes the observed treatment effect.


Estimation results. We estimate equation 42 for the first house lottery and the first cash lottery separately. We focus on first interactions with each lottery because later savings responses may be confounded by the induced intertemporal spillovers from prior lotteries and liquidity constraints. Modeling subsequent lotteries would require an assumption about how much of their liquid wealth subjects substituted from future into current periods to qualify for the earlier lottery, which neither the data nor the two-period structure of our framework can discipline. Focusing on the first cash and first house lottery avoids having to take a stand on this and keeps the mapping from Δs to $\hat{\pi}$ free of path dependence.

Figure C1 presents the estimates of the subjective winning probability $\hat{\pi}$ for the cash lottery and the house lottery. Hollow circle markers are point estimates with 95% confidence intervals; diamond markers show the true per-ticket winning probability π .⁴⁶ The true cash probability is 2.7%, but the negligible savings response implies a subjective probability close to zero.⁴⁷ For the house lottery, the true probability is 0.0096% while the estimated subjective probability is 1.3%—roughly two orders of magnitude larger than the true probability. The difference is statistically significant at $p < 0.001$. One caveat is that the reported standard errors hold calibrated inputs fixed and therefore understate the full uncertainty. This is true for any structural estimation that partially relies on calibrated parameters (DellaVigna, 2018). In Appendix C.1, we show that the

⁴⁶We compute π from Colpensiones administrative records as the number of prizes divided by total tickets in each draw. This reflects participation by all BEPS account holders, not only the experimental sample, because subjects did not know that our lotteries were restricted to a subsample.

⁴⁷The estimated subjective probability close to zero reflects the model's identification from savings responses in the first lottery of each intervention, which are negligible for the *Cash* treatment. This does not imply that all workers assign near-zero probability. As we saw in Figure 2, workers near the savings threshold increase savings, implying non-zero subjective odds for them. Since these responses are local, they do not reflect the subjective odds of the *average* subject, which is our object of interest.

Figure C1: Subjective vs True Winning Probabilities




Notes: This figure presents estimates of subjective winning probabilities ($\hat{\pi}$) compared to true ex-post winning probabilities (π) for the cash lottery and the house lottery. Hollow circle markers represent point estimates of subjective probabilities obtained from GMM estimation of equation (42), with bars indicating 95% confidence intervals. Diamond markers show the true winning probability per ticket, calculated by dividing the number of prizes by the total number of lottery tickets earned during the qualification period. The cash-lottery estimate is based on the first biweekly qualification period (April 8–21, 2019), when the partial-dose *Cash* accounts (messaged only April–June 2019) were still being treated; they are therefore included.

underestimation of small-prize lottery odds and the overestimation of large-prize ones persist under a wide range of alternative parameter calibrations.

C.1 Sensitivity Analysis for the Estimation of $\hat{\pi}$

Several calibrated parameters enter the mapping from Δs to $\hat{\pi}$, namely δ , y , and y_r . Subjects may also have incorrect beliefs about portfolio returns, R , which could affect the estimates. In Figure C2, we report a sensitivity analysis in which we vary the assumed income, the time discount factor, the payment by the retirement safety net, and the portfolio return of the retirement account over a wide grid. While the numerical point estimates shift, the qualitative conclusion is preserved across the grid: workers underestimate the odds of the low-prize lottery and overestimate the odds of the high-prize lottery.

Figure C2: Sensitivity of Subjective Odds to Calibration Choices



Notes: This figure examines the robustness of the subjective winning probability estimates from Figure C1 to alternative calibration choices. In each panel, one structural parameter is varied over a grid while all other parameters are held at their baseline values: mean monthly income of 700,000 COP, $\delta = 0.99$, $y_r = 80,000$ COP per month, $R = 1.05$, $m = 0.20$, and $\gamma = 0.0008578$. The gray dashed line shows the true per-ticket winning probability π . Panel A varies the target mean of the informal-worker income distribution used to rescale SISBEN. Panel B varies the annual discount factor. Panel C varies the monthly *Colombia Mayor* benefit entering the retirement-wealth proxy. Panel D varies the real return on retirement savings. Estimates in Panel D winsorize the worker-specific factor multiplying the subjective probability in the GMM moment at its 99th percentile within the treated group to limit the influence of outliers arising from the exponential CARA utility. As in Figure C1, the cash-lottery estimate uses the first biweekly qualification period (April 8–21, 2019); the partial-dose *Cash* accounts (messaged only April–June 2019) are therefore included.