Playing Catch-up with Health Savings

Jacob Berman, Adam Bloomfield, and Sita Slavov

November 2025

CRI WP 2025-03
Center for Retirement Initiatives Working Paper
Center for Retirement Initiatives (CRI)
McCourt School of Public Policy
Georgetown University
125 E Street NW
Washington, D.C, 20001
Email: criretirement@georgetown.edu
http://www.cri.georgetown.edu

Acknowledgements: We are grateful to Lucas Goodman, Alice Heath, Matt Litton, and seminar attendees at the Office of Tax Analysis (OTA) for helpful comments and suggestions. This research was conducted while Berman was employed at the U.S. Department of the Treasury. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors and do not necessarily reflect the views or official positions of the U.S. Department of the Treasury, the FDIC, the U.S. Government, or the Georgetown University Center for Retirement Initiatives. Any taxpayer data used in this research was kept in a secured Treasury or IRS data repository, and all results have been reviewed to ensure that no confidential information is disclosed.

Playing Catch-up with Health Savings

Jacob Berman, Adam Bloomfield, and Sita Slavov

Abstract

We use comprehensive tax data to study how saving behavior responds to the Health Savings Account (HSA) "catch-up" contribution provision, which raises HSA contribution limits for individuals aged 55 and older. Using a regression discontinuity design, we find a sharp increase in contributions among those previously near the limit and smaller increases among unconstrained savers. Induced contributions are not immediately withdrawn and do not appear to crowd out retirement savings. Responses are strongest among payroll contributors and long-term savers. However, married couples do not appear to coordinate their HSA behavior to take advantage of the complex spousal rules governing catch-up contributions. Our findings highlight how tax incentives shape HSA saving and suggest that tax-advantaged account design meaningfully affects household financial behavior.

Jacob Berman
Office of Tax Analysis, Department of Treasury
1500 Pennsylvania Avenue NW, Washington, DC 20220
lucas.goodman@treasury.gov

Adam Bloomfield
Georgetown University Center for Retirement Initiatives
125 E Street NW, Washington, DC 20001
adam.bloomfield@gmail.com

Sita Slavov
George Mason University and National Bureau of Economic Research
3351 Fairfax Drive, Arlington, VA 22201
sslavov@gmu.edu

I. Introduction

Tax-advantaged savings accounts have emerged as a preferred tool of U.S. policymakers to encourage personal savings and promote long-term financial security. Defined contribution (DC) retirement plans and Individual Retirement Accounts (IRAs) hold close to $30 trillion (ICI 2025), and nearly all individuals have at least one such account at some point in their lives (Coyne et al. 2024). In addition, Americans have accumulated more than $520 billion in 529 college savings accounts (Federal Reserve Board 2024) and $2 billion in ABLE accounts (Social Security Administration 2024). Among the various tax-preferred savings options, Health Savings Accounts (HSAs) offer the greatest tax advantages for the marginal saver. When paired with a high-deductible health plan (HDHP), contributions are exempt from payroll and income tax. Investment returns and withdrawals for health expenses are also tax-exempt, and penalty-free withdrawals for non-health expenses are available starting at age 65. Since their introduction in 2003, policymakers have periodically broadened HSA eligibility and the scope of allowable expenses, most recently through the One Big Beautiful Bill Act of 2025.[1] Despite their rapid growth and recent legislative expansions, the literature examining taxpayer behavior in response to HSAs remains limited. Prior work has leveraged data from individual employers (Leive 2022), but the relevance of these effects for the broader population is unclear.

In this paper, we provide descriptive evidence on HSA saving behavior using comprehensive tax data. We also estimate the causal effect of HSA contribution limits on HSA saving, HSA withdrawals, and DC saving. We begin by documenting trends in HSA contributions, withdrawals, and wealth accumulation. Many HSA holders use their accounts primarily for current

[1] Specifically, Public Law 119-21 authorizes pre-deductible telehealth coverage under HDHPs, allows capped HSA payments for direct primary care while treating such arrangements as compatible coverage, and deems all ACA Marketplace bronze and catastrophic plans as HSA-eligible HDHPs.

medical expenses, often contributing and withdrawing within the same year. In a balanced panel of first-time HSA contributors, we find the median balance quickly reaches zero while the average balance grows steadily--a pattern that underscores the skewed distribution of taxpayer engagement with HSAs.[2] Next, we apply a regression discontinuity (RD) design that exploits the HSA 'catch-up' provision, which creates a discontinuous increase in the HSA contribution limit at the exact age of 55.[3] We then examine how this increase in the limit affects individual HSA contributions, both overall and for different types of savers. At the household level, we examine spousal coordination in contributions when one or both spouses become eligible for catch-up contributions. Next, we estimate a marginal propensity to consume (MPC) from HSA balances by exploring the impact of the catch-up provision on withdrawals. Finally, we examine whether the increase in HSA contributions induced by the catch-up provision crowds out contributions to DC retirement plans.

Our study offers the first causal analysis of HSA policy using administrative tax data. We leverage restricted access tax records from the Internal Revenue Service (IRS) in which we can observe the universe of HSA contributions, distributions, and balances for U.S. taxpayers. In addition, the data include individual and employer characteristics, allowing for a detailed examination of how HSA saving behavior varies across demographic groups and other subpopulations. The data also include DC contributions, which help us measure spillover effects of HSA saving on other tax-advantaged savings accounts. Because the sample is large, we have sufficient statistical power to identify changes in HSA activity that would be undetectable in other contexts.

[2] To protect privacy, all medians reported in this paper are pseudo-medians calculated by taking the mean of the 10 observations around the median.

[3] Specifically, individuals who turn 55 by December 31 of a given year are eligible to contribute an extra $1,000 to their HSA, while individuals born one day later will not be eligible until the following year.

In our setting, a standard life-cycle savings model with no liquidity constraints makes three predictions: (1) individuals who are constrained by the pre-catch-up HSA contribution limit increase their contributions when they reach age 55; (2) same-year distributions from HSAs do not change (implying an MPC of zero); and (3) some savers reallocate from defined contribution (DC) plans to HSAs.

For Prediction 1, our evidence is partially consistent with the benchmark model. We find a sharp increase in HSA contributions at age 55 among individuals who contributed above 90 percent of the limit in the previous year. At the same time, we observe a more modest increase in contributions among individuals not previously constrained by the limit, a response that may reflect anchoring or heuristic-based decision-making. Rules for family plans are more complex and require spousal coordination to fully take advantage of catch-up contributions. Thus, we separately analyze the impact of the catch-up provision on married couples. We find that married couples do not appear to coordinate their contributions, which may reflect a lack of knowledge, psychological factors relating to the complexity of the rules, or intra-household bargaining.

For Prediction 2, our evidence is consistent with the benchmark model. Although the catch-up provision induces a marked increase in contributions, we find no change in same-year distributions and only a modest increase in next-year distributions. The saving induced by the discontinuity is not used immediately for near-term medical expenses; rather, taxpayer behavior is consistent with a longer-term savings motive. Our result contrasts with Leive (2022) who estimates a marginal propensity to withdraw same-year contributions of 0.85. This result also contrasts with our descriptive finding that the median individual withdraws most of their HSA savings soon after starting to contribute. This discrepancy can be explained by the fact that the compliers in our RD setting are part of the group using HSAs for longer-term savings.

For Prediction 3, our evidence is inconsistent with the benchmark model. Because HSAs have greater tax advantages than DC accounts such as 401(k)s, some individuals should reduce their DC savings to contribute more to their HSA. However, we find no evidence that becoming eligible for HSA catch-up contributions causes a change in DC contributions. Thus, the increase in the HSA contribution limit appears to increase overall saving across HSAs and DC plans.

Our paper contributes to several strands of literature. First, we add to a small body of research on the extent to which individuals respond to contribution limits and tax incentives in tax-advantaged savings vehicles. Prior work has documented responses to contribution limits in retirement accounts such as 401(k)s and IRAs. LaVecchia (2018) finds that IRA catch-up eligibility results in a substantial increase in total tax-deferred contributions among individuals who lack access to a DC savings plan. This increase occurs along both the intensive and extensive margins: there is an increase in average IRA contributions as well as the probability of making any contribution. In a related paper, Goodman (2020) finds that becoming eligible for catch-up contributions to a DC retirement account, at age 50, results in an increase in retirement contributions with no evidence of crowd-out of savings in taxable accounts. Dao and Rao (2024) find that the contribution increases associated with the retirement savings catch-up rule vary by household income level. We are the first to examine the impact of catch-up contributions within HSAs, which represent a large and growing market. A key advantage of our setting is that, relative to the DC limit, the HSA limit binds for a much higher share of contributors. This suggests the effects we measure are relevant for a broader group of taxpayers.

Second, this paper contributes to the recent literature on consumer behavior in the context of HSAs. Using data from a large employer, Leive (2022) estimates that roughly 85 percent of every additional HSA dollar is withdrawn in the same year—far above the neoclassical benchmark

of zero—challenging the idea that employees treat HSAs as long-term savings. He also finds no evidence that HSA contributions crowd out retirement savings. Likewise, Davis et al. (2023) show that most employees within their sample of 15 universities keep only modest, un-invested HSA balances, indicating that the accounts are used chiefly for short-term medical spending rather than for retirement saving Relative to these recent papers, we use much larger and more representative data and confirm that many individuals, indeed, spend a large share of their HSA savings each year. In contrast to prior research, we identify a large group of individuals who do not immediately withdraw their contributions, behaving instead in a manner consistent with longer-term savings. We also find that the additional contributions induced by the catch-up provision are rarely offset by withdrawals in the short term. A few papers in this literature have also examined interactions among health plan choice, HSA savings, and DC savings. Using data from a single employer, Leive et al. (2024) find evidence that individuals who choose a dominated health plan are less likely to take advantage of employer matching contributions for their DC plan. Friedberg et al. (2023) develop a life-cycle model incorporating HSA savings, DC savings, and liquid savings accounts and show that DC and HSA accounts are complements for workers with low HSA contribution levels (i.e., lower wage employees with liquidity constraints that include medical spending) and substitutes for those with higher contribution levels. Our empirical results, which focus on individuals with high contribution levels, suggest that HSA contributions and retirement contributions are neither complements nor substitutes.

Third, this paper contributes to the broader literature on behavioral responses to savings incentives. An important question in this literature is whether savings incentives, such as tax preferences, increase overall savings or simply shift it from unsubsidized to subsidized forms (e.g., Poterba et al. 1996; Engen et al. 1996; Hubbard and Skinner 1996; Gelber 2011; Attanasio and

DeLeire 2002; Chetty et al. 2014; Bernheim 2002). In our setting, HSAs generally have a greater tax advantage than DC accounts. Contributions to both DC plans and HSAs are exempt from income tax; however, HSA contributions made through payroll deductions are further exempt from payroll taxes. Moreover, unlike DC accounts, there is no penalty or tax on HSA withdrawals used for medical expenses. As with DC accounts, there is no penalty on withdrawals – for medical or non-medical expenses – after age 65. Thus, an increase in access to HSA savings opportunities may crowd out DC savings for many workers. However, we find that the increased HSA contribution limit at age 55 appears to increase the HSA savings of constrained individuals without crowding out (or crowding in) DC contributions. (We cannot rule out crowd-out of other kinds of savings.)

Fourth, our analysis speaks to the broader literature on the role of behavioral factors in life-cycle savings and portfolio choice (a classic paper is Shefrin and Thaler 1998; for reviews, see Ericson and Laibson 2019 and Beshears et al. 2018). We find, as Choukhmane et al. (2024) do for DC contributions and Vihriälä (2025) does for credit card payments, that couples do not jointly optimize HSA contributions at the household level, possibly reflecting psychological factors or household bargaining. We also find evidence that individuals respond to the discontinuity in the contribution limit even if they were not constrained in the years before they reached age 55. We explore the extent to which this finding can be explained by psychological factors such as anchoring (e.g., Choi et al. 2017). Finally, although theory predicts that taxpayers shift savings across accounts to optimize their after-tax return, we find that increased HSA savings at age 55 does not crowd-out or crowd-in defined contribution (DC) retirement savings. We provide a discussion of whether this failure to reallocate savings can be rationalized in a standard life-cycle model.

This study advances the existing literature in several important ways. Although previous research has explored responses to catch-up contributions in retirement accounts, little is known about how individuals respond to similar provisions in HSAs. Additionally, our use of comprehensive IRS tax records provides a level of granularity and coverage rarely available in studies of individual saving behavior. The data allow us to capture the universe of HSA holders and precisely measure both contributions and distributions. This contrasts with prior studies that often rely on survey data or administrative records from a single employer— sources that are less representative and may overlook key segments of the population. By leveraging a natural experiment and high-quality administrative data, we provide novel evidence on the behavioral responses to the HSA saving opportunities.

II. Background

Health Savings Accounts (HSAs) were established in 2003 with the aim of encouraging individuals to save for medical expenses. Individuals must be enrolled in a High-Deductible Health Plan (HDHP) to contribute to an HSA. Because of this requirement, the tax subsidies for HSAs reward individuals for choosing HDHPs, which in turn are meant to incentivize cost-conscious health care consumption by increasing out-of-pocket spending before insurance coverage begins.

HSAs offer unique tax advantages that distinguish them from other savings vehicles. Contributions to an HSA are exempt from federal payroll and income tax, investment returns grow tax-free, and withdrawals are tax-free when used for qualified medical expenses. This "triple tax advantage" makes HSAs the most tax-favored savings vehicles available to individuals in the United States. Unlike funds in Flexible Spending Arrangements (FSAs), which generally expire after one year, HSA balances never expire. Assets in an HSA can be invested and rolled over

indefinitely, allowing long-term accumulation. Unlike DC contributions made by employees, HSA contributions made through payroll deductions are exempt from payroll taxes. Additionally, once an individual reaches age 65, HSA withdrawals for non-medical expenses are permitted without penalty (though they are subject to ordinary income tax), further enhancing the account's appeal as a supplemental retirement savings vehicle.

The IRS sets annual HSA contribution limits based on whether the linked HDHP covers an individual or family. To encourage additional savings among older individuals who tend to have higher health expenses, individuals aged 55 and older are eligible for "catch-up" contributions. In recent years, this catch-up provision has allowed for an additional $1,000 contribution beyond the standard limit. Couples with family HDHP coverage can jointly contribute up to the family limit. They can either split their contribution across two separate accounts or contribute to a single account in one spouse's name. (HSAs do not allow joint ownership.) However, catch-up contributions must be made to an account held by the individual who is eligible for additional contributions. Thus, for a couple using a single HSA in one spouse's name, taking full advantage of catch-up contributions requires opening and contributing to a second account in the other spouse's name.[4]

Since their inception, HSAs have experienced significant growth. Our data show that the number of individuals with an HSA exceeds 25 million, with assets exceeding $110 billion as of 2023.

III. Data, Summary Statistics, and Descriptive Results:

a. Data

[4] There is no similar rule for distributions, which may be used to cover the health expenses of the account holder, their spouse, or their dependents regardless of who made the contribution initially.

Our analysis uses administrative tax data from the IRS, covering the universe of HSAs. This microdata provides detailed information on HSA contributions, distributions, balances, retirement savings, and individual characteristics such as age and date of birth. The richness and granularity of these data enable us to precisely measure how individuals adjust saving behavior in response to the HSA catch-up provision. Specifically, we leverage data from four key IRS forms: Form W-2, Form 5498-SA, Form 1099-SA, and Form 8889. Form W-2 provides employer-reported information on annual compensation, including DC retirement savings and pre-tax HSA contributions made through payroll deductions.[5] This information allows us to distinguish between employee contributions made through workplace payroll systems and those made independently, outside the workplace. Form 5498-SA is filed by HSA custodians and reports total annual contributions to HSAs, regardless of the contribution channel. The form also reports HSA balances on the last day of the year as well as "late" contributions.[6] Form 1099-SA reports distributions from HSAs, enabling us to analyze when taxpayers use their HSA funds. Form 8889 is filed by individual taxpayers with HSAs and reports whether they are enrolled in a self-only plan or a family plan. We observe taxpayers' sex and birth dates by linking Social Security numbers to Social Security Administration records.

Our sample includes individuals aged 54.5 to 55.5 as of December 31 of each calendar year from 2010 to 2022. The restriction ensures that the analysis focuses on a narrow window around the policy cutoff to limit the influence of nonlinear, age-related confounders. Additionally, it allows

[5] Our measure of DC retirement savings includes traditional and Roth employee contributions. We do not observe employer DC contributions. In contrast, our measure of HSA contributions includes payroll contributions from both the employer and the employee.

[6] Late contributions are contributions made between January 1 and April 15 of year *t+1* that count towards the limit in year *t*. They are reported in Box 3 of Form 5498-SA.

us to stack multiple years together without double-counting any individuals. Nominal dollar amounts are converted to constant 2023 dollars using the chained Consumer Price Index.

b. Summary Statistics

Table 1: Characteristics of Individual HSA Savers age 54.5-55.5 over years 2010-2022

Characteristic	**HSA Contribution > 0**	**HSA Contribution > 90% of Current Year Limit**
Male (share)	0.59	0.65
Married (share)	0.70	0.77
Family Plan (share)	0.60	0.64
Any W2 (share)	0.86	0.90
Any W2 with HSA payroll Contribution (share)	0.76	0.74
AGI (median)	131,344	207,069
HSA Contribution (mean)	3,539	7,119
HSA Balance (mean)	5,045	11,679
Any HSA Withdrawal (share)	0.80	0.83
HSA Withdrawal (mean)	2,493	4,232
Any DC (share)	0.68	0.74
DC Contribution (mean)	8,009	14,008
Observations	2,956,647	774,927

Notes: This table reports the sample of individuals that file Forms 1040 and Form 8889 and are of age 54.5 to 55.5, as of December 31 in each tax year (2010-2022). All dollar values indexed to 2023 using Chained CPI.

Table 1 presents summary statistics. The unit of observation is the individual taxpayer. Our sample construction focuses on individuals approaching the HSA catch-up contribution eligibility threshold at age 55, ensuring that the regression discontinuity design captures behavioral responses precisely at the age cutoff. We report statistics for individuals who made HSA contributions, and for a subset whose contributions exceeded 90 percent of the statutory limit.[7]

[7] As a simplifying assumption, we calculate the limit without respect to contributions for the taxpayer's spouse. We defer a discussion of spousal contributions to section V.e.

Men predominate among HSA savers in this age group, and their share rises further in the high-contribution group. The increase in male representation as contribution levels rise suggests potential sex differences in savings decisions, which we explore further in the heterogeneity analysis.

Withdrawals are higher in the high contribution group, though the growth in withdrawals is somewhat less pronounced than that of contributions. The average HSA balance for high contributors is $11,679, more than double the average balance for all HSA contributors. Defined contribution (DC) retirement saving is common in this sample: 68 percent of all contributors, and 74 percent of high contributors, make DC contributions. Average DC contributions increase from $8,009 to $14,008 across the sample groups, suggesting that heavy HSA savers are also high savers in other tax-advantaged accounts.

c. Descriptive Results

Figure 1: HSA Distribution to Contribution Ratios (2021 Contributors)

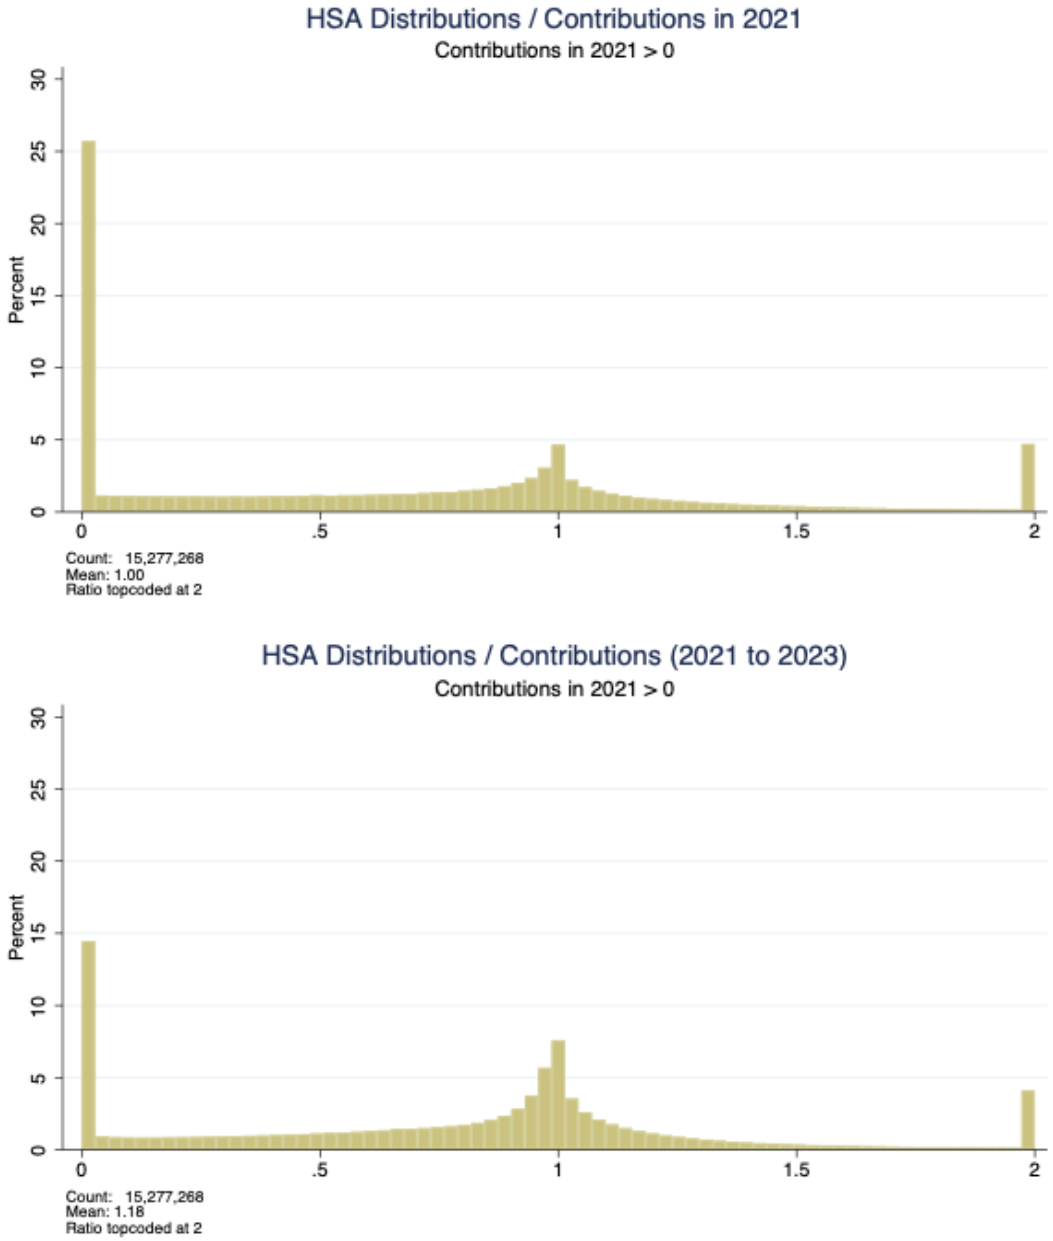

Notes: This figure reports the distribution of the ratio of HSA distributions to contributions for a sample of taxpayers with positive HSA contributions in 2021. The top panel reports the distribution to contributions ratio for 2021 alone, while the bottom panel reports the cumulative distribution to contribution ratio over 2021-2023.

Figure 1 presents histograms of the distribution-to-contribution ratio for individuals with positive HSA contributions in 2021. Each histogram is top-coded at a ratio of 2 for readability. The histograms show whether HSAs are used mainly for spending or for saving.

The top panel displays the ratio of distributions to contributions for 2021 alone. The large mass at zero (over 25 percent of the sample) indicates that more than one in four individuals with positive contributions in 2021 took no distributions at all that year—clear evidence of net saving behavior in the short term. Though the modal ratio is 0, there is a second sizeable concentration of mass around a ratio of 1 (peaking just below 5 percent of individuals). These individuals tend to spend roughly what they contribute, consistent with pass-through use of the HSA for tax-advantaged spending within a year.

The bottom panel repeats this analysis over a three-year horizon, summing both contributions and distributions across 2021–2023 for individuals who had positive contributions in 2021. The distribution flattens and the mass at zero falls to just below 15 percent, indicating that many of those who did not withdraw their contributions in 2021 eventually did so over the following two years. This behavior is consistent with delayed spending rather than long-term savings. The mass around the ratio of 1 grows substantially, to between 5 and 10 percent, and the surrounding bins (ratios just above and below 1) are also notably denser than in the prior one-year figure. The shift suggests that over a slightly longer horizon, a greater share of HSA users spend approximately what they contribute. The mass at the top-coded ratio of 2 remains relatively stable between the two panels.

Overall, the one- and three-year distributions suggest that most HSA contributors spend some or most of their contributions, while a minority consistently delay withdrawals. This finding reinforces the dual role of HSAs as both short-term liquidity tools and long-term savings vehicles.

Figure 2: HSA Contributions, Distributions, and Balances by Year (2015 First-Time Contributors)

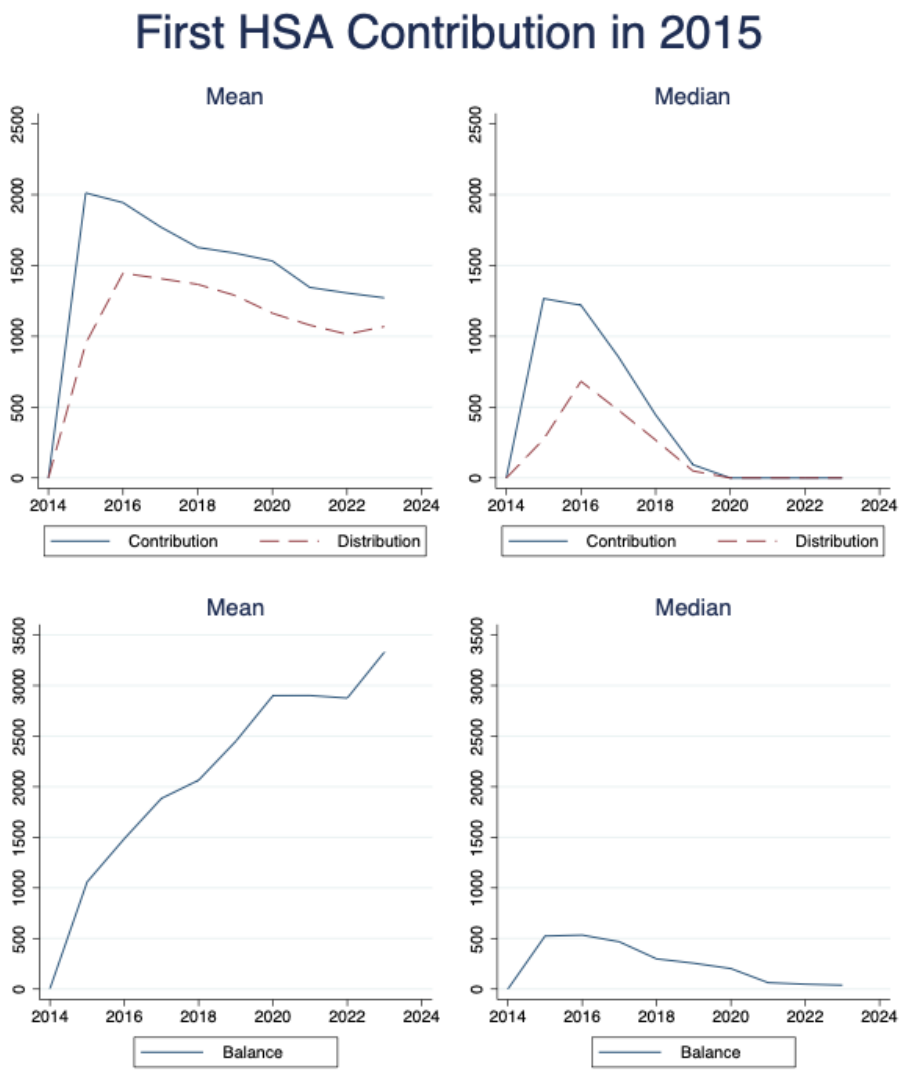

Notes: This figure tracks mean contributions and distributions (top left), median contributions and distributions (top right), mean balances (bottom left), and median balances (bottom right) for the cohort of 2.4 million individuals who made their first observed HSA contributions in 2015.

We next consider a cohort of individuals who made their first HSA contribution in 2015. Because first-time HSA participation is rising over time, we selected 2015 to balance sample size and the length of the time series. The sample includes taxpayers who become newly eligible to contribute because of a change in health insurance, receive a new employer contribution, or newly

decide to contribute despite prior eligibility. We construct a balanced panel so that individuals who made no HSA contributions or distributions after 2015 are recorded as zero rather than missing.[8]

The top two panels of Figure 2 present annual contributions and distributions from 2014 to 2023. The top left panel displays means, and the top right panel shows medians. Mean contributions peak in 2015, the first year of HSA activity, and decline steadily in subsequent years before leveling out. The front-loaded mean contribution intensity reflects attenuation in participation. In contrast, mean distributions reach their peak in the second year of HSA use (2016), before gradually leveling off. The divergence between 2015 and 2016—declining contributions alongside rising distributions—suggests that new savers begin to use their HSAs more actively for spending after an initial accumulation period. Therefore, same-year withdrawals offer a limited picture of spending dynamics.

Medians depicted in the top right panel are lower, and median distributions also lag behind median contributions. Median contributions spike in 2015 but fall to zero by 2019, remaining there through the end of the period. Median distributions follow a similar pattern, peaking in 2016 and declining to zero by 2019. The typical HSA user in this cohort becomes inactive within a few years of opening the account.

Together, these panels show that for the 2015 cohort of first-time HSA contributors, engagement is highest in the first year of saving. Distributions lag contributions, consistent with a pattern of saving first and spending later. While average behavior reflects ongoing use with declining intensity after the initial period, median behavior implies that most users disengage from active HSA saving and spending within several years.

[8] We exclude a trivial share of taxpayers who ever have a "non-normal" distribution. Non-normal distributions include those related to corrections of excessive contributions, prohibited transactions, disability, or death.

The bottom two panels of Figure 2 show the evolution of mean (bottom left) and median (bottom right) HSA balances among our cohort. Mean account balances rise consistently, indicating that a subset of HSA users continues to contribute and invest funds. The steady balance growth suggests persistent saving behavior among longer-term users. In contrast, the median balance is around $500 in 2015, the first year of contributions. It remains relatively steady in 2016 before declining sharply in subsequent years, reaching nearly zero by 2023. This pattern suggests that the typical account holder exits the HSA system, with balances drawn down or left idle.

The divergence between mean and median balance trajectories underscores the skewed nature of long-term HSA saving: while a minority of users steadily build wealth in their accounts, the majority disengage soon after entry, leaving a near-zero balance within several years.

IV. Methods

To identify the causal effect of the HSA catch-up contribution provision, we employ a regression discontinuity (RD) design that exploits the sharp increase in the allowable HSA contribution limit at age 55. This policy-induced discontinuity in contribution limits provides a natural experiment to examine how individuals adjust their saving behavior in response to the additional incentive. The RD design is well-suited to this setting as eligibility for the catch-up provision changes discretely at a known and easily enforced cutoff.

The running variable in our RD design is the individual's exact age on December 31, measured in days. We define the running variable as the number of days relative to the individual's 55th birthday, evaluated on December 31 (negative before 55, positive after). The primary estimating equation is a stacked RD with varying linear trends in age:

$$Y_i = \alpha + \beta D_i + \gamma Age_i + \delta(D_i Age_i) + \epsilon_i \qquad (1)$$

where Y_i represents the outcome of interest for individual i. Our primary outcomes include total annual HSA contributions, total annual HSA distributions, and defined contribution (DC) retirement plan contributions. D_i is an indicator equal to 1 if the individual is eligible for catch-up contributions (i.e., age $\geq$ 55) and 0 otherwise. The coefficient of interest, β, captures the local average treatment effect of the catch-up provision among compliers, and ϵ_i is an idiosyncratic error term. We collapse all variables at the date of birth level and weight the regression by the number of individuals within a cell. To account for correlated shocks among individuals with the same birthday, we cluster standard errors at the day-by-month level. To estimate an MPC, we additionally implement an instrumental variable regression using the date of birth discontinuity to instrument for contributions. In the instrumental

Figure 3: Regression Discontinuity Estimate of the Effect of the HSA Contribution Limit Increase on HSA Contributions

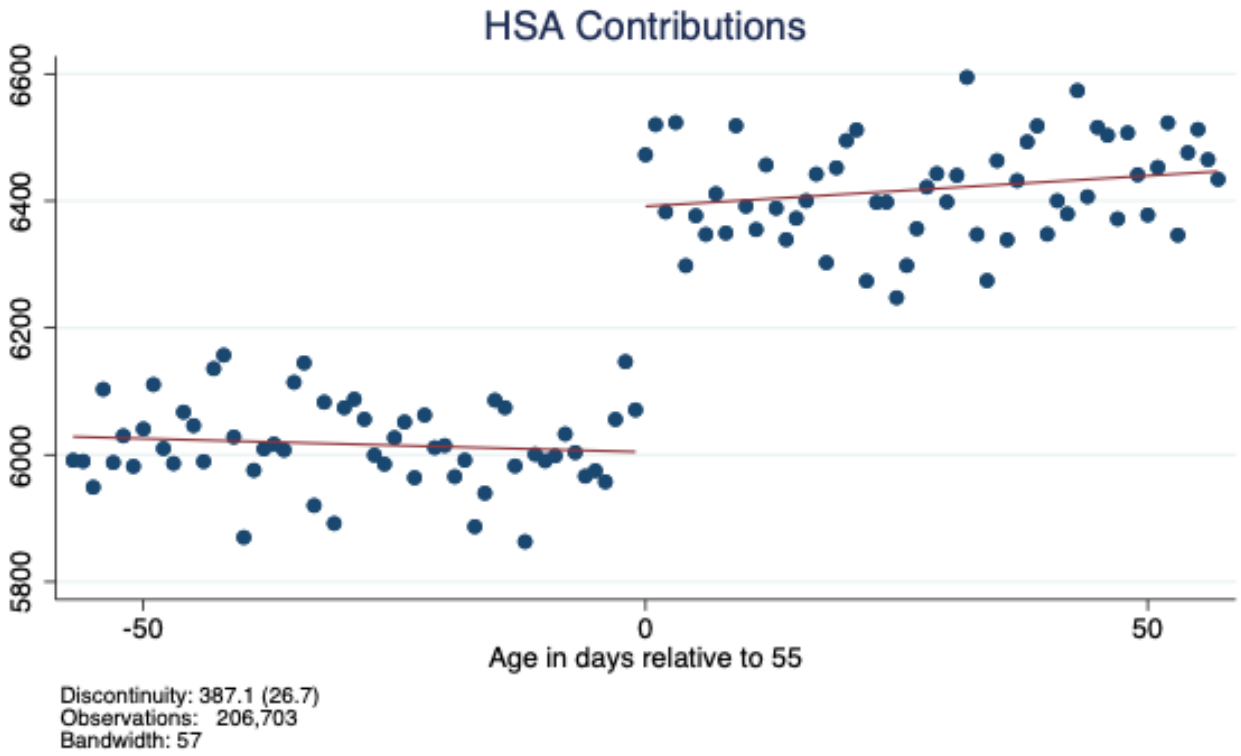

Notes: Figure depicts the regression discontinuity estimate of the HSA catch-up contribution provision based on equation (1). Data are from 2010-2022 and the sample is restricted to individuals whose prior-year HSA contributions exceeded 90 percent of the statutory limit.

variables model, equation (1) serves as the first stage. We compute the optimal bandwidth for equation (1) based on the nonparametric methods in Calonico, Cattaneo, and Titiunik (2014). For consistency, we use the same bandwidth for the heterogeneity analysis and disaggregations, as well as to obtain the IV estimate of the MPC.

The coefficient of interest, β, captures the local average treatment effect of the catch-up provision among compliers if two identification assumptions are satisfied. First, taxpayers cannot precisely manipulate their date of birth around the cutoff. Second, no potential confounders are also changing discontinuously at the cutoff.

V. Regression Discontinuity Results

a. Main finding: effect on HSA contributions

Figure 3 displays the regression discontinuity (RD) estimate of the effect of the HSA catch-up contribution provision on annual HSA contributions. The sample includes individuals whose prior-year contributions were above 90 percent of the applicable limit, ensuring that the analysis focuses on savers who were plausibly constrained. The figure plots average HSA contributions as a function of individuals' exact age (in days) relative to December 31 of the tax year, within a 57-day window around the age-55 threshold. Solid lines depict local linear regressions fitted separately on either side of the cutoff.

The results show a sharp increase in HSA contributions at the age-55 threshold. The discontinuity is both precisely estimated and economically significant: individuals just above age 55 contribute an average of $387.1 more annually than those just below the eligibility age. This

response represents a 6.4 percent increase relative to the mean contribution below age 55 and amounts to 38.7 percent of the $1,000 increase in the statutory contribution limit. While this average response accounts for a substantial share of the $1,000 limit increase, it still implies that most savers do not fully exploit the additional tax-advantaged saving capacity afforded by the catch-up provision.[9]

b. Validating the Regression Discontinuity Design

Our ability to measure a causal effect depends on the assumption that taxpayers are unable to manipulate their date of birth near the threshold. Potential threats could arise from strategic retiming of births around December 31 or fraudulent changes to birth records. To test for manipulation of the running variable, Figure A2 plots the counts of taxpayers within each date of birth cell. Consistent with prior work using tax data (Goodman 2020), we find that the distribution of dates of birth is not perfectly smooth around December 31. There are modest shifts in counts surrounding winter holidays as births near Christmas appear to be delayed by several days. While it is implausible that this manipulation occurs for the purpose of maximizing HSA contributions, we perform an additional check by implementing two placebo tests. First, we re-estimate equation (1) in our primary sample using HSA contributions in the tax year before individuals turn 55 (that is, tax years 2009 to 2021). Second, we construct a new sample of taxpayers turning 54 and examine current year HSA contributions around age 54 (tax years 2010 to 2022). Figures A3 and A4 plot the respective results and show no evidence of a discontinuity. These results suggest that there are no unobserved confounders associated with date of birth that are discontinuously influencing HSA contributions.

[9] The outcome in Figure 3 is HSA contribution that are made *for* a tax year. Figure A1 shows that redefining the outcome as contributions *in* a tax year does not affect the result.

An additional threat to identification is discontinuous changes in other policies around the age-55 cutoff. We are aware of one such policy in the context of DC plan withdrawals. The "Rule of 55" lets taxpayers take penalty-free distributions from the 401(k) or 403(b) plans associated with their most recent employer if they separate from employment in the calendar year they turn 55. Plan sponsors are not obligated to offer this feature, and any distributions allowed by the plan remain subject to ordinary income tax. If taxpayers in this situation are incentivized to convert their DC distributions into HSA contributions, then the Rule of 55 may complicate our interpretation of the policy treatment. This scenario is unlikely for several reasons. First, the rule is not widely known, and it is only available in narrow circumstances and at the discretion of the employer. Second, individuals who can take advantage of this flexibility are likely to be liquidity constrained and uninterested in tax incentives for additional savings. Similarly, individuals who are separating from employment may be losing the HDHP coverage that makes them eligible to contribute to an HSA. Finally, individuals can already convert retirement balances into HSA balances at any age by rolling over DC accounts into an IRA and making a qualified funding distribution to the HSA. However, to test the impact of the Rule of 55 empirically, Figure A5 plots retirement plan distributions for our primary sample using information on Form 1099-R. We see no change in retirement distributions, suggesting the Rule of 55 is not impacting HSA contributions.

c. Heterogeneity in main findings

Table 2 disaggregates the results from Figure 3 by three key characteristics—sex, HDHP coverage type, and prior-year HSA withdrawal history—to examine heterogeneity in treatment effects across subpopulations. For each subgroup, we report the mean HSA contribution among

individuals below age 55, the RD estimate of the catch-up provision based on equation (1), the standard error of the estimate, the percent change relative to the pre-55 mean, and the corresponding sample size.

The estimated discontinuity is larger for men than women; however, the level of contributions for men is consistently higher both before and after the threshold. Thus, the catch-up contribution provision induces similar percentage increases for both sexes.

Table 2: Heterogeneity analysis by sex, coverage type, and prior withdrawal

Subsample	**Mean below 55**	**Discontinuity**	**Standard Error**	**Percent Change**	**Sample Size**
Male	6,394	425.9***	(31.6)	6.66%	136,963
Female	5,269	323.2***	(37.1)	6.13%	69,740
Self-Only	3,846	326.7***	(31.0)	8.49%	71,182
Family	7,144	427.8***	(31.4)	5.99%	135,521
Prior Distribution	6,065	352.0***	(30.1)	5.80%	174,497
No Prior Distribution	5,478	579.1***	(54.0)	10.57%	32,206

Notes: Regression discontinuity estimates by group of catch-up contribution provision based on equation (1). Data are from 2010-2022 and the sample is restricted to individuals whose prior-year contributions exceeded 90 percent of the statutory limit. *** p<0.01, ** p<0.05, * p<0.10.

When disaggregating by plan type, we find that contribution levels among those with family HDHP coverage are roughly double those with self-only coverage, reflecting the larger statutory limits. The estimated discontinuity is larger in absolute terms for family coverage, though the percent increase is greater among those with self-only coverage.

We also disaggregate the estimates based on whether individuals took any HSA distributions in the prior year. This split helps differentiate between two types of savers: those who typically use their annual HSA savings for current medical expenses—benefiting from spending pre-tax earnings—and those who forgo withdrawals and are more likely to treat HSAs as a long-term savings vehicle, benefiting additionally from tax-advantaged growth. We find significant effects for both groups, but the largest response (in both absolute and relative terms) occurs among

longer-term savers. Although this group exhibits the largest discontinuity of any subsample, the magnitude remains well below the $1,000 increase in the statutory limit. Even among this high-saving group, many do not fully exploit the increased limit.

d. Decompositions of main findings

Table 3: Main result decomposed by contribution method and time

Outcome	Mean below 55	Discontinuity	Share of Total Discontinuity	Standard Error
Payroll	4,211	275.2***	71.1%	(31.8)
Non-Payroll	1,805	111.9***	28.9%	(29.1)
Late Contribution	287.4	52.5***	13.6%	(9.9)
Timely Contribution	5,729	334.6***	86.4%	(26.2)

Notes: Table shows regression discontinuity estimates by group of catch-up contribution provision based on equation (1). Data are from 2010-2022, and sample is restricted to individuals whose prior year contributions exceeded 90 percent of the statutory limit. *** p<0.01, ** p<0.05, * p<0.10.

We next decompose the first-stage response by contribution method and timing. The most common method for making HSA contributions is through payroll deductions, which is optimal for tax purposes because it is the only way to exclude contributions from payroll taxes.[10] Alternatively, taxpayers can make non-payroll deposits directly to the HSA custodian—similar to IRA contributions. Contributions may be made during or after the tax year, with the deduction claimed when the taxpayer files their tax return.

Theory offers ambiguous predictions about whether payroll or non-payroll contributions will respond more strongly to the catch-up provision. On the one hand, prior research suggests that employees exhibit substantial inertia in their decisions to participate in and contribute to employer-

[10] Employers are incentivized to offer this method not only because it allows them to save on the employer share of payroll taxes. Self-employed workers cannot deduct HSA contributions for the purposes of the self-employment tax. We ignore the impact of payroll contributions on Social Security earnings which may impact Social Security benefits.

sponsored retirement plans. This inertia underlies a large literature—originating with Madrian and Shea (2001)—showing that automatic enrollment and default options in employer-sponsored retirement plans significantly increase participation and influence investment decisions. This research shows that workers who are automatically enrolled in a retirement plan tend to stick with default contribution and investment options, implying that "passive" savers tend to be less responsive to policy changes. In contrast, savers and investors making active choices have been shown to be much more responsive to tax policy changes than passive savers and much less responsive to automatic savings regimes. Using Danish data, Chetty et al. (2014) find that "active" savers are significantly more responsive to tax subsidies than passive savers whereas passive savers are more responsive to automatic contribution regimes.[11] In the HSA context, these prior findings suggest that many workers who contribute via an automatic mechanism—payroll deduction—may behave as passive savers. Due to inertia, these individuals may be less likely to adjust their savings behavior in response to the increased contribution limit at age 55. On the other hand, workers making payroll contributions may be more financially sophisticated than average. Moreover, employers or co-workers may provide informational nudges to workers making payroll contributions – e.g., reminding them of increased limits. Thus, payroll contributions may alternatively be more responsive to the catch-up limit. The net effect is therefore theoretically ambiguous.

Table 3 shows contributions made by payroll deduction rise by $275.2 at the cutoff, while non-payroll contributions rise by $111.9. The increase in payroll contributions represents 71 percent of the overall increase induced by the catch-up provision; that share is roughly in line with

[11] The authors estimate that about 85 percent of individuals are passive savers while 15 percent are active savers.

mean pre-55 payroll contributions as a share of pre-55 total contributions.[12] Thus, the catch-up provision does not disproportionately affect payroll or non-payroll contributions. But the presence of substantial non-payroll contributions raises a different issue. Given that most of these non-payroll contributors are W-2 employees, it is puzzling that they do not choose to contribute via payroll, which provides a larger tax advantage by virtue of the payroll tax deduction.

Table 3 also shows that both late and timely contributions rise in response to the catch-up provision. The increase in late contributions represents 14 percent of the increase in total contributions. In comparison, late contributions represent only 5 percent of pre-55 contributions. Thus, the catch-up provision seems to disproportionately affect late contributions compared to timely contributions. That finding is consistent with delayed awareness of the catch-up provision. If some individuals become aware of the catch-up contribution after the end of the current year, they may be induced to make late contributions.

e. Impact of Catch-Up Contributions on Previously Unconstrained Individuals

Table 4: Regression Discontinuity Estimate of Contribution Limit Increase on Previously Unconstrained Individuals

Sample	Mean below 55	Discontinuity	Percent Change	Standard Error	Bandwidth	Sample Size
0% to 50%	1,797	42.2***	2.35%	(11.1)	64	405,655
50% to 75%	3,832	88.2***	2.30%	(23.9)	60	129,722
75% to 90%	4,904	138.7***	2.83%	(41.9)	58	56,027

Notes: Table shows regression discontinuity estimates of the catch-up contribution provision based on equation (1). Data are from 2010-2022, and sample is restricted to individuals whose prior year contributions were 75-90 percent (top row), 50-75 percent (middle row), and 0-50 percent (bottom row) of the statutory limit. *** p<0.01, ** p<0.05, * p<0.10.

[12] Because each contribution type (e.g., payroll vs. non-payroll, or late vs. timely) is mutually exclusive and collectively exhaustive, the sum of the discontinuities across rows approximates the total first-stage effect. Reporting each as a share of that total provides a clear decomposition of the overall response. This approach avoids misleading scale comparisons across subgroups and helps highlight the relative contribution of each channel to the aggregate policy effect.

While individuals already contributing near the pre–catch-up limit may readily increase savings, those who contributed well below the limit may respond differently. To examine this lower-contribution group, we re-run our first stage in three new samples defined by prior-year contributions: 75–90 percent of the limit, 50–75 percent of the limit, and 0–50 percent of the limit. Table 4 shows that all three of these unconstrained subgroups respond to the policy change by increasing their contributions. Although the absolute increase in contributions is smaller for individuals with low prior-year contributions, the relative percentage change is consistent across groups, ranging between 2.3 and 2.8 percent. The presence of a significant discontinuity even among previously unconstrained savers suggests that the catch-up provision may affect savings behavior through psychological mechanisms. An alternative explanation is that the jump reflects compositional changes in the sample. Each year, some previously unconstrained individuals may experience health, preference, or employment shocks that make HSA saving more attractive, pushing them into the constrained group. Likewise, some previously constrained individuals may experience shocks that make HSA saving less attractive, causing them to become unconstrained. By construction, Table 4 includes individuals transitioning from the unconstrained group to the constrained group, but not those moving in the opposite direction. Thus, part of the observed discontinuity may reflect the behavior of newly constrained individuals rather than solely behavioral responses among the truly unconstrained.

To further examine the role of potential compositional changes, we decompose a taxpayer's contribution for a given year into a direct and indirect effect. Contributions (C_i) for individual i can be decomposed as:

$$C_i = \max(C_i - l_i, 0) + \min(C_i, l_i)$$

where, l_i is individual i's pre-catch-up contribution limit. The first term only applies to individuals who would be constrained by the pre-catch-up limit and captures the portion of their contribution that exceeds this limit. We refer to the policy's impact on this term as its *direct effect*. The second term captures the *indirect effect.* For unconstrained individuals the term is equal to their contribution and for constrained individuals the term equals the limit.

Table 5 reports the decomposition within our RD approach. We estimate that the direct effect of the catch-up provision is around $121 and the indirect effect is around $51. These results suggest that the policy has a substantial (though smaller) effect on saving even on contributions that would not be affected by the pre-catch-up limit.[13]

Table 5: Direct versus Indirect Effect Based on Current Year Contribution

Outcome	**Mean below 55**	**Discontinuity**	**Standard Error**	**Share of Total Discontinuity**
Total Contributions	3,435	171.4***	(16.5)	100.0%
Direct Effect	58.4	120.6***	(2.3)	70.4%
Indirect Effect	3,377	50.8***	(15.9)	29.6%

Notes: Sample includes filers of Form 8889 from 2010 to 2022 who made current-year contributions and range from age 54.5 to 55.5. No restriction is placed on prior- year contributions. Bandwidth is 51 days. Sample size = 822,989 individuals. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

These results reinforce the notion that compositional changes alone cannot explain the observed effects among previously unconstrained individuals. A more likely explanation is the presence of psychological mechanisms, such as anchoring, framing, or heuristic-driven decision-making. For example, the higher contribution limit may serve as a salient reference point or a reminder that one could—or "should"—be saving more, prompting an upward adjustment in contributions even among individuals who were not previously constrained. Alternatively, some individuals may anchor their saving decisions to a proportion of the maximum limit (e.g., aiming

[13] The fact that the mean below 55 for the direct effect is slightly above zero shows there some non-compliance with the limit.

to save "half" or "three-quarters" of the limit), and adjust contributions upward when the limit increases. The gradient in discontinuity magnitudes across the three rows of Table 4 is consistent with this interpretation.

f. Do spouses optimize their contributions around the HSA catch-up contribution limit?

Table 6: Spousal Coordination of Catch-Up HSA Contributions

Subsample	**Mean below 55**	**Discontinuity**	**Standard Error**	**Percent Change**	**Sample Size**
Ignoring Spousal Rules	7,185	440.0***	(33.5)	6.12%	103.817
Accounting for Spousal Rules	4,327	254.8***	(44.3)	5.89%	176,988
Spousal Rules Bind	286	0.1	(16.2)	0.03%	73,171
Spouse w/ No Prior Contribution	17.6	4.4	(6.2)	25.00%	67,744

Notes: Table presents regression discontinuity estimates of the catch-up contribution provision based on equation (1). Data are from 2010-2022. See text for details of sample construction. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 6 reports RD estimates exploring whether spouses coordinate their contributions to take full advantage of catch-up contributions. As described above, when a married couple is covered by a family HDHP, their joint HSA contribution limit is the sum of the family contribution limit and each spouse's individual catch-up contribution. Contributions may be made to a single account held in one spouse's name or to two separate accounts held individually by each spouse. When either of the two spouses covered by the family HDHP turns age of 55, that spouse is eligible to make catch-up contributions. However, the catch-up contributions must be made to an account held in the name of the spouse who is eligible for them. For example, if a couple contributes solely to an account in the younger spouse's name, they would need to open a new HSA in the older spouse's name to take advantage of the catch-up rule.

The top row of Table 6 shows the impact of the catch-up provision on contributions for individuals who are married and enrolled in a family HDHP. The sample is restricted to individuals whose prior year contribution was at least 90 percent of the family contribution limit. This row

ignores the pooling of family contributions by implicitly assuming that spousal contributions are zero. For example, consider a married couple where one spouse contributed $300 and the other $7,000. Suppose the family limit is $7,300. The spouse contributing $300 would be excluded from the sample in the top row even though the family is constrained by the contribution limit. The spouse contributing $7,000, however, would be included in the sample. The estimated discontinuity at age 55 is $440, in line with the main results from Figure 3.

The second row acknowledges spousal rules by including all individuals who are part of a constrained household. That is, it includes individuals whose contributions are within 90 percent of the family limit minus the spouse's contributions. In the same $300/$7,000 example, both spouses are included in the sample. Under this correction, the discontinuity falls to $254.8, suggesting weaker responsiveness when low-contributing spouses are counted. The attenuation in the estimated effect points to a lack of coordination across spouses.

To explore this result further, the third row isolates those whose individual contributions are not within 90 percent of the family limit but are part of a household that is constrained. In other words, these individuals are excluded from Row 1 but included in Row 2. Here, the observed discontinuity is a precisely measured zero. The results indicate there is no detectable behavioral response among low-contributing spouses in constrained households. The bottom row restricts the sample to individuals in constrained households who made no HSA contributions in the previous year, possibly because they do not have an HSA in their name. We similarly find no effect.

Taken together, the results in Table 6 reveal that while the catch-up contribution provision has a large impact on those whose individual contributions place them close to the family limit, it has no impact on their lower- or non-contributing spouses. Even when lower-contributing spouses are positioned to increase their household's HSA limit due to their age, there is no measurable

adjustment in savings behavior. This lack of responsiveness may reflect poor understanding of the complex spousal rules or household bargaining considerations (e.g., taking full advantage of the catch-up limit may require one spouse to transfer control of money to the other).

g. Does the HSA catch-up contribution limit affect withdrawals?

Figure 4: Effect of the HSA Contribution Limit Increase on Withdrawals

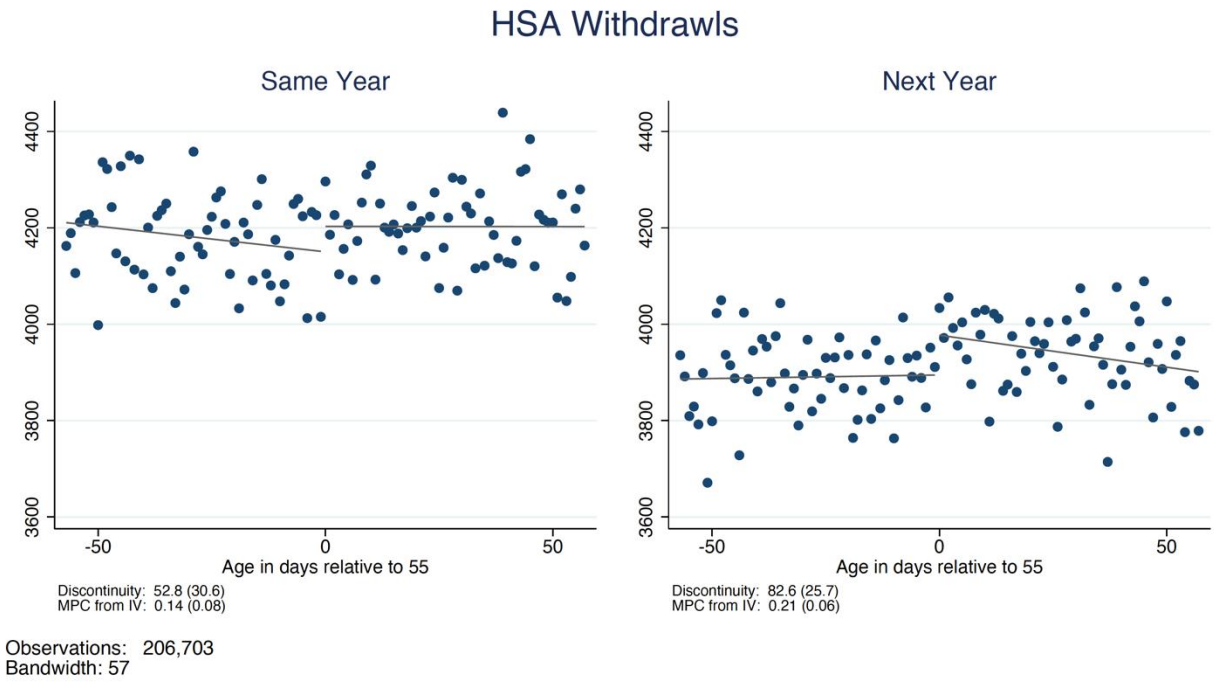

Notes: Figure depicts regression discontinuity estimate based on equation (1). Data are from 2010-2022, and sample is restricted to individuals whose prior year contributions exceeded 90 percent of the statutory limit.

Figure 4 reports regression discontinuity (RD) estimates of HSA withdrawals at the catch-up eligibility threshold of age 55, focusing on individuals whose prior-year contributions were at or above 90 percent of the statutory limit. Although Figure 3 showed a clear discontinuity in contributions at age 55, we do not observe a statistically significant discontinuity in same year withdrawals. Using our instrumental variables specification, we reject an MPC greater than 0.29 for same-year withdrawals. Individuals who increase their savings due to the catch-up provision

are not simultaneously increasing their withdrawals, at least in the short term. If HSAs for this group were being used primarily as pass-through accounts—where individuals contribute pre-tax dollars only to withdraw them shortly thereafter—one might expect to see a parallel jump in distributions at the eligibility threshold.

In the year following age-55 eligibility (t+1), we estimate a statistically significant, albeit still modest, increase of $82.6 in withdrawals, implying an MPC out of contributions of 0.21. While a portion of their additional saving is eventually drawn down, the effect remains small relative to the increase in contributions due to the catch-up threshold, indicating that most of the additional contributions reflect genuine savings behavior rather than immediate tax-advantaged medical spending. Our MPC is notably lower than the 0.85 estimated by Leive (2022). One explanation is that the compliers in our setting are more likely to be long-term savers. Our primary sample has higher income and larger baseline contributions than Leive's. In addition, Leive's identification relies on exogenous variation in employer HSA contributions, whereas ours does not. Employees may treat employer contributions as "extra" or less earmarked for future needs and therefore may be more willing to spend them quickly. In contrast, their own contributions may be part of a deliberate saving plan. Such mental accounting could lead to a higher MPC in Leive's setting.

We next examine responses by historical withdrawal behavior. Table 7 groups individuals by the share of contributions withdrawn over the prior three years: (1) none, (2) less than 50%, (3) 50–90%, and (4) more than 90%. Among groups (1), (2), and (3), there is no statistically significant increase in distributions in year t or t+1. Those who withdrew less than 50 percent of their contributions over the previous three years exhibit a significant decline in distributions of $109.95 in year t. Overall, these results indicate that even among those who withdrew large shares of their

HSA contributions in recent years, there is no immediate increase in withdrawals at age 55. Among individuals who historically withdrew almost all their annual contributions, there is a statistically significant increase in distributions of $139 in year *t* and $116 in year *t+1*. Thus, unlike the other groups, these individuals appear to adjust their HSA spending behavior alongside their increased saving. However, even for the group that is historically inclined to spend as intensely as they save, only about a third of their induced catch-up contributions are distributed for additional spending.

Table 8 presents regression discontinuity (RD) estimates assessing whether the HSA catch-up contribution provision at age 55 leads to adjustments in defined contribution (DC) retirement savings. This analysis investigates potential substitution effects, wherein individuals who increase their HSA contributions at the threshold may offset this increase by reducing their contributions to DC plans (e.g., 401(k)s). The analysis focuses on those who contributed at or above 90 percent of the prior year's HSA contribution limit. The top row

Table 7: Withdrawals Discontinuity by Three-Year Ratio of Withdrawals to Contributions

		Distributions in T			
Subsample	**Sample Size**	**Mean below 55**	**Discontinuity**	**SE**	**Percent Change**
Ratio = Zero	21,552	595	-2.0	(48.8)	-0.34%
0 < Ratio < 0.5	49,519	2,713	-74.4	(58.2)	-2.74%
0.5 < Ratio < 0.9	67,620	5,022	33.3	(43.7)	0.66%
Ratio > 0.9	68,012	5,537	139.1***	(49.5)	2.51%
		Distributions in T+1			
Subsample	**Sample Size**	**Mean below 55**	**Discontinuity**	**SE**	**Percent Change**
Ratio = Zero	21,552	714	35.9	(61.8)	5.03%
0 < Ratio < 0.5	49,519	2,632	65.3	(60.6)	2.48%
0.5 < Ratio < 0.9	67,620	4,577	38.3	(45.7)	0.84%
Ratio > 0.9	68,012	5,118	115.7**	(51.6)	2.26%

Notes: Table presents regression discontinuity estimates of HSA withdrawals, in period *t* and *t+1*, induced by the HSA catch-up contribution provision based on equation (1). Data are from 2010-2023, and samples are stratified by withdrawal intensity over the prior three years. *** p < 0.01, ** p < 0.05, * p < 0.10.

h. Does the HSA catch-up contribution limit affect DC retirement savings?

Table 8: Regression Discontinuity Estimates of the Effect of the HSA Contribution Limit Increase on DC Contributions

Subsample	Mean below 55	Discontinuity	Standard Error	Percent Change	Sample Size
Primary Sample	13,927	18.6	(107.1)	0.13%	206,703
Below DC Limit	8,817	85.1	(57.0)	0.97%	149,533

Notes: Table 8 presents regression discontinuity estimates for DC savings induced by the HSA catch-up contribution provision based on equation (1). Data are from 2010-2022, and samples are restricted to individuals whose prior year HSA contributions exceeded 90 percent of the statutory limit. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

of the table shows that the estimated discontinuity in DC contributions at age 55 is a statistically insignificant $18.6, suggesting that individuals who increase their HSA contributions upon becoming eligible for the catch-up provision do not systematically reduce their DC savings in response. The lack of a significant downward adjustment in DC contributions suggests that the increase in HSA contributions represents a net increase in total tax-advantaged savings, rather than a reallocation across savings accounts characterized by unequal degrees of tax advantage.

The sample used in the first row of the table includes individuals who are constrained by the DC contribution limit and those who are not. However, high-saving individuals who are constrained by both the HSA and DC contribution limits may be at a corner solution with respect to tax-preferred saving; these individuals are not likely to reduce DC contributions upon becoming eligible for HSA catch-up contributions. The second row of the table excludes these high savers by restricting the sample to individuals who were not contributing at or above 90 percent of the DC maximum in the prior year ($t-1$). That is, it provides results for people who are constrained by the HSA contribution limit but not the DC contribution limit – and therefore are most likely to reduce their DC contributions when they become eligible for HSA catch-up contributions, which are more tax-advantaged than DC saving. For this subset of savers, the estimated discontinuity in DC contributions at age 55 is a statistically insignificant $85.1. Thus, even among individuals who

were not maximizing their DC contributions before age 55, there is no meaningful reduction in DC savings in response to the increased HSA contribution limit. The absence of such an adjustment further suggests that savers do not actively reallocate and optimize their tax-advantaged portfolio between these two types of savings vehicles in response to the HSA catch-up provision.

Overall, the results indicate that increasing HSA savings at age 55 does not lead to a systematic decrease in DC retirement savings. This behavior may reflect the existence of discontinuous differences in the after-tax return across savings vehicles, which give rise to kinks in individuals' budget constraints. For example, consider an individual who contributes just enough to a DC plan to exhaust the employer match: they are at a corner with respect to the DC plan's incentives. In other words, the last dollar contributed to the DC plan earned a higher return than the next dollar contributed will. Thus, it may not be optimal to reduce DC contributions at the margin when HSA catch-up contributions become available despite their relatively higher tax advantage. Alternatively, many savers may be following default contribution settings for their DC plans, making them less likely to actively adjust their 401(k) contributions in response to the new HSA opportunity. Finally, the observed behavior may suggest that individuals do not treat HSA saving as a relatively close but more tax-advantaged substitute for DC saving. In the last two cases, it is possible that individuals fail to optimize the allocation of their long-term savings.

VI. **Conclusion**

This paper provides empirical evidence on HSA-related consumer behavior and presents the first causal analysis of how individuals respond to the catch-up contribution provision, leveraging rich administrative tax data covering the universe of HSA holders. We find that becoming eligible for the age-55 catch-up provision leads to a sharp increase in HSA contributions

among individuals previously constrained by the statutory limit, consistent with standard economic theory. We also document smaller but significant behavioral responses among individuals who were not previously constrained, potentially driven by psychological factors.

Our analysis reveals considerable heterogeneity in responses to the catch-up provision, with stronger effects among those contributing via payroll deduction and those with no recent history of HSA withdrawals. Despite large increases in savings, we find no immediate rise in HSA distributions and only modest increases in distributions the following year, indicating that much of the induced saving reflects longer-term saving. We find no evidence that increased HSA savings crowds out retirement contributions to defined contribution (DC) plans, even among individuals with room to adjust. These findings suggest that catch-up contributions expand total tax-preferred saving rather than reshuffling it across vehicles.

We also explore how complex contribution rules—particularly for married individuals—affect household saving behavior. Our evidence shows that spouses fail to optimize their HSA contributions jointly, particularly when taking full advantage of catch-up provisions would require shifting contributions between the two spouses' accounts. This pattern may reflect institutional complexity, behavioral frictions, or household bargaining dynamics.

Several limitations should be noted when interpreting the results. First, our analysis focuses on HSAs and DC retirement accounts. We do not consider other vehicles such as 529 plans, FSAs, or regular taxable savings accounts. This limits our ability to assess whether HSA saving substitutes for other forms of saving outside of retirement accounts. Second, HSA distributions do not necessarily correspond to contemporaneous health spending, as individuals may withdraw funds to reimburse themselves for expenses incurred at any time after the account was established. Finally, we cannot observe whether distributions were used for a particular type of health service

(such as preventive care versus emergency treatment), which constrains our ability to assess the implications for health care utilization. Future work that merges HSA tax data with health claims data would be valuable for identifying the timing and nature of medical spending associated with HSA use.

Taken together, our findings underscore how tax incentives shape saving behavior in tax-advantaged accounts. As HSAs continue to grow in importance, understanding their role within the broader savings landscape will be critical for designing effective savings policy.

References

Attanasio, Orazio P. and Thomas DeLeire. 2002. "The Effect of Individual Retirement Accounts on Household Consumption and National Saving." *The Economic Journal* 112(481): 504-538.

Bernheim, B. Douglas. 2002. "Taxation and Saving." In *Handbook of Public Economics*, Volume 3 (ed. Alan Auerbach and Martin Feldstein). Elsevier: 1173-1249.

Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian. 2018. Behavioral Household Finance. In *Handbook of Behavioral Economics: Applications and Foundations*, Volume 1 (eds. B. Douglas Bernheim, Stefano DellaVigna, David Laibson). North-Holland: 177-276.

Board of Governors of the Federal Reserve System (US), Households; Net Worth, Level [HNOA529P], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/HNOA529P.

Chetty, Raj, John N. Friedman, Søren Leth-Petersen, Torben Heien Nielsen, and Tore Olsen. 2014. "Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark." *Quarterly Journal of Economics* 129(3): 1141-1220.

Choi, James J., Emily Haisley, Jennifer Kurkoski, and Cade Massey. 2017. "Small Cues Change Savings Choices." *Journal of Economic Behavior & Organization* 142: 378-395.

Choukhmane, Taha, Lucas Goodman, and Cormac O'Dea. 2025. "Efficiency in Household Decision Making: Evidence from the Retirement Savings of US Couples." *American Economic Review* 115(5): 1485-1519.

Dao, Ngoc and Manita Rao. 2024. "Does the Catch-up contribution policy improve retirement preparedness?" *Work, Aging, and Retirement* December 2024.

Davis, Brent, Adam Leive, and Andrew Gellert. 2023. "Fungibility in workplace benefits choices: Evidence from Health Savings Accounts." *TIAA Institute Research Paper Series* No. 198.

Engen, Eric M., William G. Gale, and John Karl Scholz. 1996. "The Illusory Effects of Saving Incentives on Saving." *Journal of Economic Perspectives* 10 (4): 113-138.

Ericson, Keith Marzilli and David Laibson. 2019. Intertemporal Choice. In *Handbook of Behavioral Economics: Applications and Foundations*, Volume 1 (eds. B. Douglas Bernheim, Stefano DellaVigna, David Laibson). North-Holland: 177-276.

Friedberg, Leora, Jaeki Jang, Adam Leive, and Eric Young. 2023. "Health Savings Accounts and Life-Cycle Saving: Implications for Retirement Preparedness" Working Paper.

Gelber, Alexander M. 2011. "How Do 401(k)s Affect Saving? Evidence from Changes in 401(k) Eligibility." *American Economic Journal: Economic Policy* 3 (4): 103-22.

Goodman, Lucas. 2020. “Catching up or crowding out? The crowd-out effects of catch-up retirement contributions on non-retirement saving.” *Journal of Public Economics* 188: 104221.

Hubbard, R. Glenn and Jonathan S. Skinner. 1996. “Assessing the Effectiveness of Saving Incentives.” *Journal of Economic Perspectives* 10(4): 73-90.

Leive, Adam. 2022. “Health insurance design meets saving incentives: consumer responses to complex contracts.” *American Economic Journal: Applied Economics* 14(2): 200-227.

Leive, Adam, Leora Friedberg, and Brent David. 2024. “Links Between Puzzles in Household Finance: Evidence from Employee Benefit Choices.” Working Paper.

Poterba, James M., Steven F. Venti, and David A. Wise. 1996. “How Retirement Saving Programs Increase Saving.” *Journal of Economic Perspectives* 10 (4): 91-112.

Shefrin, Hersh M and Richard H. Thaler. 1988. “The Behavioral Life-Cycle Hypothesis.” *Economic Inquiry* 26(4): 609-643.

Social Security Administration. 2024. “ABLE Act: 10 Years of Progress for People with Disabilities.” Social Security Matters Blog, March 27, 2024. https://blog.ssa.gov/able-act-10-years-of-progress-for-people-with-disabilities.

Vihriälä, Erkki. 2025. “Intra-Household Frictions, Anchoring, and the Credit Card Debt Puzzle.” *The Review of Economics and Statistics* 107 (2): 510-522.

Appendix

Figure A1: Robustness Check for Contribution in a Tax Year

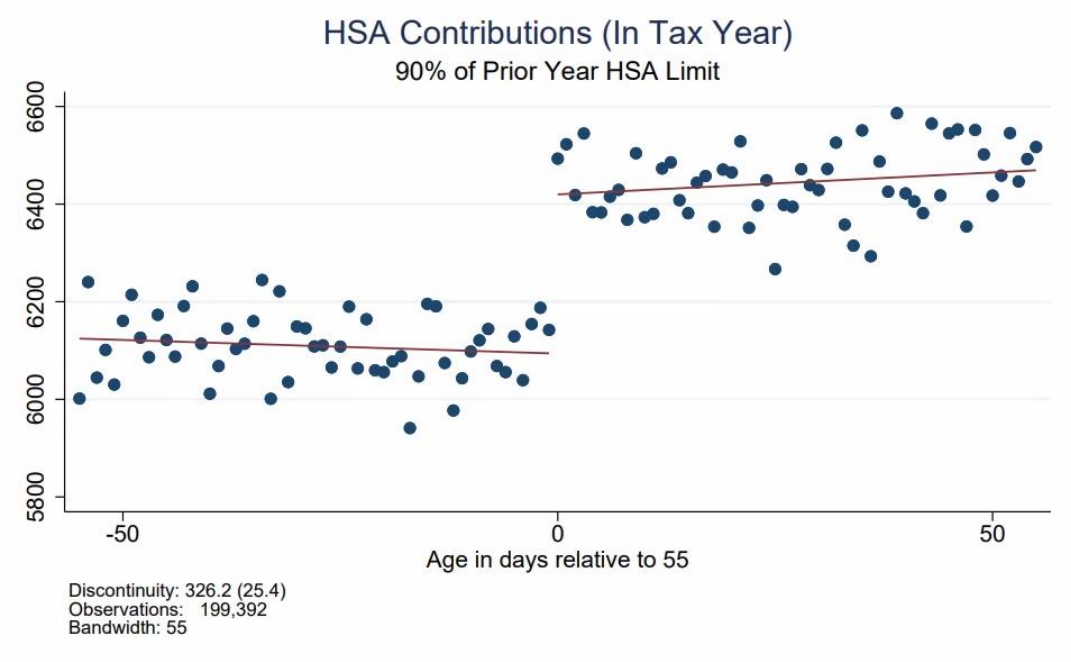

Figure A2: Counts of Birthdates within each Day

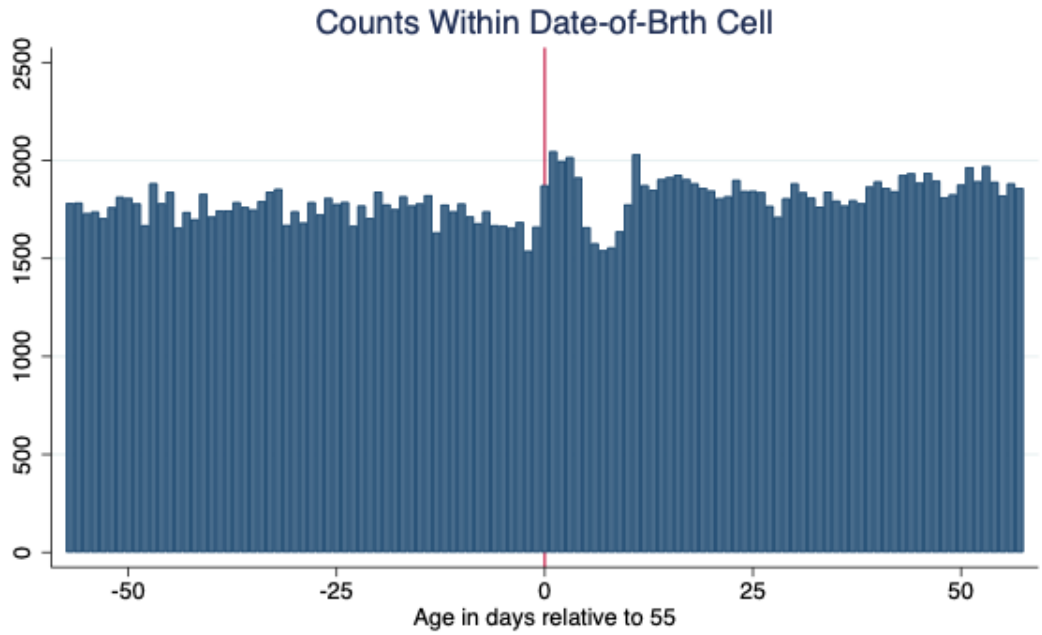

Figure A3: Placebo Test RD Estimate of HSA contributions in Tax Year before Primary Sample Individuals turn Age 55 (2009-2021).

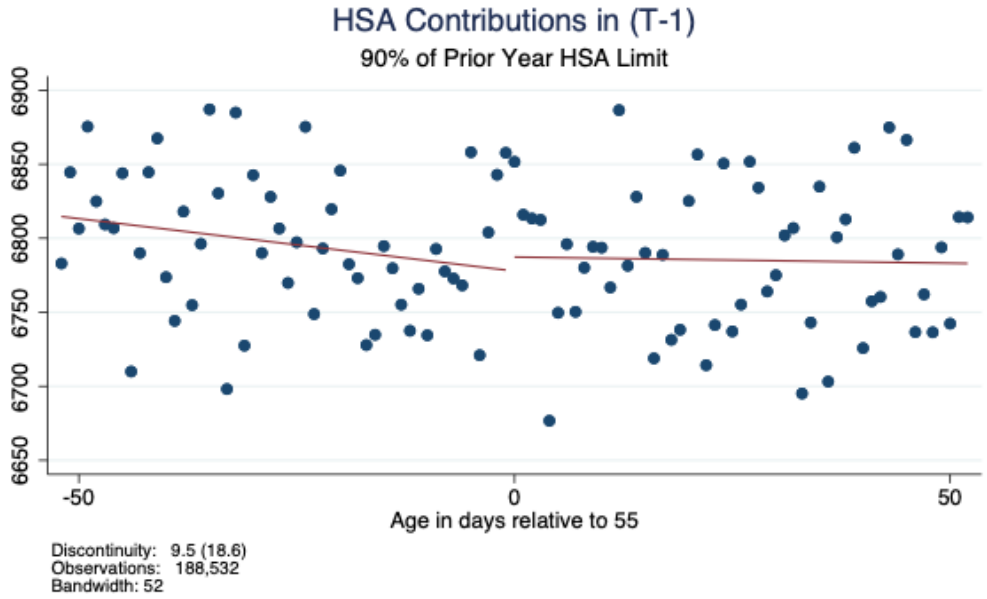

Figure A4: Placebo Test RD Estimate of HSA contributions for Individuals turning Age 54 (2010-2022).

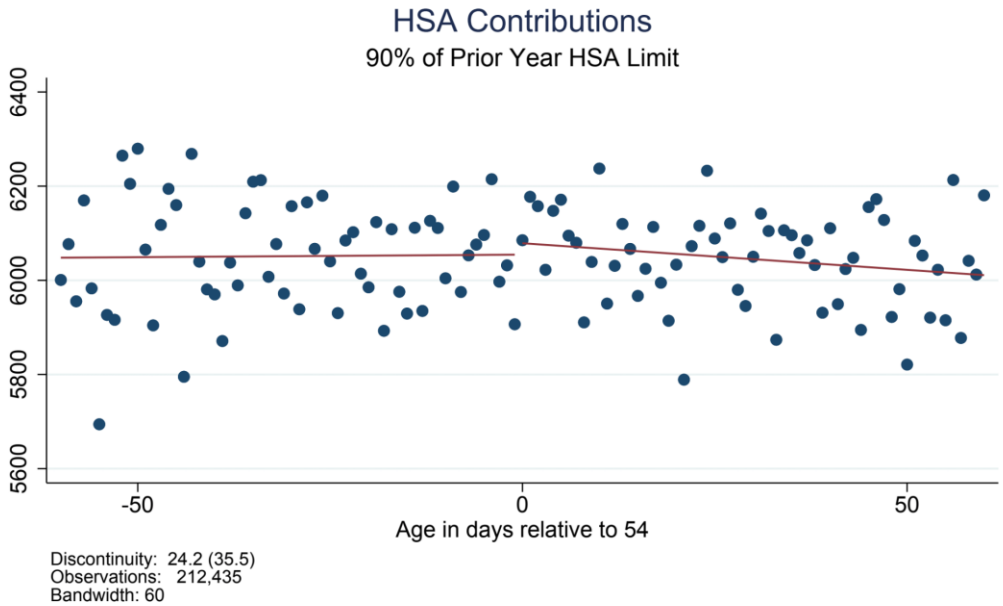

Figure A5: RD Test of the "Rule of 55" on DC Retirement Plan Distributions