Temporary and Permanent Effects of Withdrawal Penalties on Retirement Savings Accounts*

Gopi Shah Goda
Stanford University
and NBER

Damon Jones
University of Chicago

Shanthi Ramnath
Federal Reserve Bank
of Chicago

October 2022

Abstract

We investigate the impact of the early withdrawal penalty on Individual Retirement Accounts (IRAs) withdrawals by examining behavior in a short window before and after the age when the penalty is lifted. We find a large, sudden increase in withdrawals after the penalty's expiration of 3-3.5 times the baseline level, with no evidence of anticipatory behavior. After one month, average withdrawals decline to and persist at approximately double the baseline. We find that the short-run increase is more pronounced when liquidity constraints are more likely to have been binding. Finally, we explore the implications of our results for policies that adjust the age at which early penalties expire or policies that temporarily remove such penalties.

JEL Classification: H24; D14; J32

Keywords: retirement savings accounts, withdrawals, distributions, penalty

*Email: gopi@stanford.edu; damonjones@uchicago.edu; sramnath@frbchi.org. We would like to thank Konhee Chang, Paula Gablenz, Rolando Hernandez Gomez, Yalun Su and Andre Vasilyev for outstanding research assistance. Goda and Jones acknowledge generous support from the U.S. Social Security Administration through grant #RRC08098400-08 to the National Bureau of Economic Research as part of the SSA Retirement Research Consortium and by the TIAA Institute and the Pension Research Council/Boettner Center of the Wharton School at the University of Pennsylvania. The authors also thank Adrien Auclert, Alex Gelber, Greg Kaplan, David Laibson, Hannes Malmberg and seminar participants at Stanford, Consumer Financial Protection Bureau, Williams College, University of Virginia, American University and the U.S. Treasury for helpful comments. The findings and conclusions expressed are solely those of the authors and do not represent the views of SSA, the US Department of the Treasury, the Federal Reserve Bank of Chicago, the Federal Reserve System, the TIAA Institute, the Pension Research Council, or the NBER.

1 Introduction

Contributions to Individual Retirement Accounts (IRAs) and employer-sponsored defined contribution plans (e.g. 401(k) plans) receive preferential tax treatment, allowing retirement wealth to grow at a faster rate. In turn, these savers face a penalty on withdrawals made before a specified age. Though there are exceptions to the penalty, this feature makes tax-preferred retirement accounts less liquid than a standard savings account. For example, contributions into Traditional IRAs are not subject to income tax until they are withdrawn; however, withdrawals prior to age $59\frac{1}{2}$ are subject to a 10 percent early withdrawal penalty in addition to any income tax liability. If the individual claims a "hardship exemption" as defined by the Internal Revenue Code, the early withdrawal penalty is waived.

In addition, a number of recent laws allow access to these retirement accounts penalty-free, under certain conditions. For example, individuals impacted by Hurricane Katrina could take penalty-free withdrawals. Similarly, the Coronavirus Aid, Relief, and Economic Security (CARES) Act allowed limited penalty-free withdrawals for those who experienced adverse financial consequences from COVID-19. Such measures are generally used to reduce hardship, provide insurance, and facilitate consumption smoothing.

In this paper, we examine how the increase in liquidity once one ages out of the penalty affects withdrawals from IRAs in the weeks and months surrounding the penalty's expiration at age $59\frac{1}{2}$ and as far out as two years. Our analysis uses tax records from the full sample of individuals born between July 1, 1941 and July 1, 1951 from tax years 1999 through 2013, approximately one third of whom have positive IRA balances in the calendar year they turn $57\frac{1}{2}$. For our estimation strategy, we develop a flexible, parametric method that exploits differences in one's access to penalty-free withdrawals within a calendar year, stemming from variation in date of birth.

We decompose the withdrawal response into three key parts: a short-run increase in withdrawals immediately following penalty expiration, a longer-term increase in average withdrawals, and a potential anticipatory decrease in withdrawals prior to penalty expiration.

We find that in the short-run, withdrawals from IRAs increase by a factor of 3-3.5 relative to a pre-$59\frac{1}{2}$ baseline calculated in the 30-day period prior to the removal of the early-withdrawal penalty. The increase is sudden and occurs immediately at age $59\frac{1}{2}$ when the penalty is removed. Within one month following the penalty expiration, there is a decline in average withdrawals to a persistent, longer-term level that is approximately double the baseline. We find no robust evidence of a temporary reduction in withdrawals just prior to the penalty expiration, suggesting that households are not strategically delaying withdrawals until just after the penalty expires.

We explore heterogeneity in our results by various characteristics. We find the observed patterns are present even for those in the highest quartile of either adjusted gross income (AGI) or baseline IRA wealth. The increase in the short run is more pronounced among those who turned $59\frac{1}{2}$ during the Great Recession and those with a recent history of receiving unemployment insurance (UI), where liquidity constraints are more likely to have been binding. By contrast, those with a recent history of receiving disability insurance payments do not experience a sharp increase, likely because penalties are waived for early distributions made on account of a permanent disability.[1]

Finally, we consider the welfare implications of our results for policies that affect the early withdrawal penalty. In particular, we consider an adjustment to the age at which the penalty expires, either to an earlier or later age in the neighborhood of age $59\frac{1}{2}$. We also consider a penalty holiday, which lasts for a temporary period of time and applies to all account holders below age $59\frac{1}{2}$. We show how the parameters we estimate serve as inputs into the calculation of the marginal value of public funds (MVPF) associated with these policies, as in (Hendren and Sprung-Keyser, 2020). In a baseline model, both policies involve a transfer to those who make early withdrawals and a loss in revenue from penalties that would have been collected on those early withdrawals. This cost is offset by an earlier realization of taxable income via withdrawals in the current period—rather than at the later date when the funds would

[1] It is important to note that these comparisons are among those with positive IRA account balances at baseline, a sample that is relatively higher income and wealthier than the average tax filer.

have been withdrawn in the absence of the policy—which results in an overall increase in tax revenue. In our simplest case, we find a positive MVPF from shifting the age at which the early withdrawal penalty expires a year earlier, and a potentially infinite MVPF for a year-long penalty holiday among account holders who are significantly younger than $59\frac{1}{2}$. However, we note that these results might differ under alternative sets of assumptions where the envelope theorem may not apply, including frictions or biases that lead to undersaving for retirement or significant borrowing constraints among account holders, or if there are multiplier effects from spending out of retirement accounts.

Our paper builds on an existing literature that looks at how withdrawals from retirement savings accounts change over the lifecycle. Past studies have found that pre-retirement withdrawals coincide with adverse shocks (Amromin and Smith, 2003; Argento, Bryant and Sabelhaus, 2015) and likely indicate a high marginal utility of consumption (Coyne, Fadlon and Porzio, 2020). Withdrawals can also respond to specific features of tax-preferred retirement accounts. For example, Agarwal, Pan and Qian (2020) find that both withdrawal behavior and consumption are influenced by early withdrawal penalties in Singapore, using detailed administrative records on financial transactions.

Although U.S. retirement accounts are found to be more liquid with lower penalties and several exceptions for early withdrawals compared to other developed countries (Beshears et al., 2015), a number of U.S.-specific studies find that withdrawals still respond to changes in account liquidity. In particular, Deryugina, Kawano and Levitt (2018) show that in the presence of a temporary exemption on the early withdrawal penalty, victims of Hurricane Katrina increased withdrawals from retirement accounts. Similarly, other studies have found that loosening the restriction on the mandatory age for taking withdrawals influenced withdrawal behavior (Poterba, Venti and Wise, 2013; Brown, Poterba and Richardson, 2017; Mortenson, Schramm and Whitten, 2019; Stuart and Bryant, 2020; Horneff, Maurer and Mitchell, 2021). Stuart and Bryant (2020) look at both early withdrawal penalties and the mandatory distribution age to examine IRA withdrawal behavior over the full lifecycle and estimate a dynamic model that allows the authors to identify welfare-improving adjustments to both age thresholds and penalty rates.

We contribute to this literature on responses to retirement savings accounts in several ways. First, by exploiting differences in date of birth, our approach allows us to estimate immediate, temporary spikes along with longer-term permanent increases in withdrawals following penalty expiration, unlike other studies that examine withdrawal behavior at lower frequencies. The high-frequency nature of the response is especially important when considering the policies that change incentives within a given year. Second, we find evidence that households with greater liquidity constraints tend to have a greater short-run response to penalty removal, consistent with other evidence on shocks and early withdrawals. However, we also show that in some cases, exemptions such as those for disability insurance recipients, mitigate this pattern. Finally, our method of leveraging variation in date of birth to convert lower frequency data into a higher frequency event study can be applied more generally to study other settings with age-based policies using annual data, such as the onset of Social Security eligibility. These techniques are similar to, but distinct from, techniques that exploit differences in the distribution of temperature each year to identify the effect of daily temperatures on outcomes (Deschênes and Greenstone, 2011; Deryugina and Hsiang, 2014).

2 Background and Data

2.1 Background

Tax-preferred retirement savings plans allow individuals to contribute funds annually, up to a set maximum, to accounts facilitated by an employer (e.g., 401(k) plans) or to IRAs. Contributions are either made with pre-tax assets and taxed when withdrawn (Traditional IRAs or 401(k) plans) or made with after-tax assets and exempt from taxes when withdrawn (Roth IRAs or Roth 401(k) plans). In either case, the tax-preferred treatment affords savers a higher return on savings through the deferral of taxation on earnings.[2]

These accounts impose penalties against withdrawing prior to retirement that vary by account type. Traditional IRAs in particular allow early withdrawals for any reason, but withdrawals prior to attaining $59\frac{1}{2}$ are subject to a 10 percent penalty. Exceptions to the penalty are made in the event of death or permanent disability, for first-time homebuyers, education expenses, health insurance premiums while unemployed, and unreimbursed medical expenses.[3] While the penalty on early withdrawals may encourage savings by incentivizing money to remain in an account, it may also discourage savings by decreasing liquidity, making the net effect on savings ambiguous.

Withdrawals prior to age $59\frac{1}{2}$ from 401(k) plans can be made only in the event of a hardship, or an immediate and heavy financial need. Certain expenses are deemed to be immediate and heavy, including qualifying medical expenses, the purchase or repair of a principal residence, and burial or funeral expenses. These early withdrawals are subject to a 10 percent penalty, with some exceptions (e.g., upon the death or disability of the account holder). Once the account holder turns $59\frac{1}{2}$, all penalties and restrictions are lifted.

2.2 Data

Our data come from the population of tax and information returns collected by the Internal Revenue Service (IRS). We use supplementary information provided by the Social Security Administration (SSA) on date of birth, gender, and date of death to restrict our sample to individuals born between July 1, 1941 and July 1, 1951, during tax years 1999 through 2013, who are alive in the year in which they turn $57\frac{1}{2}$. This sample restriction ensures that our data contain two tax years before and after each individual turns $59\frac{1}{2}$. Our dataset contains information on household income (Form 1040), withdrawals from IRAs and

[2] For example, assets held in a Traditional IRA for 30 years, with a 3% annual return and an income tax rate of 30%, paid upon withdrawal, will result in a retirement balance that is 30% larger than an identical amount saved in a taxable account, which must pay income taxes on the principal upon deposit and on interest earned in each interim year.

[3] Since Roth IRA contributions are made on an after-tax basis, withdrawals of the basis—but not the earnings—can be done without penalty.

employer-sponsored retirement plans (Form 1099-R), contributions to and account balances of IRAs (Form 5498), and tax amounts on early withdrawals (Form 5329). Appendix A provides additional details regarding sample construction.

We focus on withdrawals from IRAs due to some important data limitations. First, there is no tax form at the individual level that reports account balances for defined contribution plans such as 401(k) plans, making it difficult to select a sample of individuals who have a 401(k) balance. Second, withdrawals from defined contribution plans are indistinguishable from defined benefit payments on Form 1099-R. By contrast, IRA withdrawals can be separately identified due to a checkbox on the 1099-R tax form. As described previously, the penalties differ somewhat for 401(k) plans and IRAs; therefore, generalizing our results to other types of accounts should be done with caution. However, IRAs may be more typical, particularly at ages close to $59\frac{1}{2}$, since many individuals roll over employer-sponsored accounts into IRAs before retirement (Goodman et al., 2019).

Our sample contains individuals who have a positive fair market value (FMV) reported on Form 5498 in at least one IRA in the year they turn $57\frac{1}{2}$. Our data represent 12,445,149 tax filers, or 36 percent of the population who attains age $57\frac{1}{2}$ in our analysis period. Just under half (49 percent) of our sample is male and 74 percent file a joint return in the year they turn $57\frac{1}{2}$. The average AGI in our sample at age $57\frac{1}{2}$ is \$134,842. This value is relatively high because we focus on both those with assets in IRA accounts and those at older ages.[4]

Unconditionally, the average annual withdrawal increases from \$1,827 in the calendar year someone turns $57\frac{1}{2}$ to \$1,967, \$3,188, \$3,903 and \$4,081 in the calendar years they turn $58\frac{1}{2}$, $59\frac{1}{2}$, $60\frac{1}{2}$ and $61\frac{1}{2}$, respectively. The fraction of our sample that takes withdrawals from their IRA is 7-8 percent prior to the year in which the individual turns age $59\frac{1}{2}$, then increases to 16-17 percent during the two years following age $59\frac{1}{2}$. The amount withdrawn conditional on taking a distribution is approximately \$24,000 annually, and does not change

[4] We perform a validation exercise using data from the Health and Retirement Study (HRS) which yields a sample with similar observable characteristics. Appendix A provides more details regarding the sample selection procedure and the validation exercise.

markedly around the age $59\frac{1}{2}$ threshold, suggesting that the increase in average withdrawals occurring at $59\frac{1}{2}$ may be on the extensive margin.

3 Empirical Methods

While tax data provide a rich source of information on IRA withdrawals, they present a challenge in identifying short-run responses to the removal of the penalty at age $59\frac{1}{2}$. Withdrawals reported on 1099-R forms are aggregated over the full calendar year for a given account. In order to estimate withdrawal patterns at a sub-annual frequency, our empirical strategy leverages variation in exposure to the penalty-free withdrawal period driven by date of birth. For example, consider two people, one born on June 30, 1950 and another born on July 1, 1950. According to IRS rules, the former turns $59\frac{1}{2}$ on December 30, 2009, while the latter turns $59\frac{1}{2}$ on January 1st, 2010. Differences in their annual withdrawals in 2009 can be related to the fact that one has experienced two days of penalty free withdrawals, while the other faces the penalty the entire year.

To illustrate how this variation helps identify responses to the early withdrawal penalty, Figure 1a displays annual withdrawals for calendar years in which people attain ages $57\frac{1}{2}$, $58\frac{1}{2}$, $59\frac{1}{2}$, $60\frac{1}{2}$ and $61\frac{1}{2}$, grouped by their exposure to penalty-free withdrawals. We group people by their month of birth. Thus, those with one month of penalty-free withdrawals represent those attaining age $59\frac{1}{2}$ in December, those with two months of penalty-free withdrawals reach age $59\frac{1}{2}$ in November, and so forth.

As shown in Figure 1a, the level and change in average annual withdrawals between the years in which age $57\frac{1}{2}$ and age $58\frac{1}{2}$ occur are the same across different months of birth. However, the change in annual withdrawals begins to diverge in the year in which people turn $59\frac{1}{2}$. In particular, those who have more months of penalty-free withdrawal in the year in which they turn $59\frac{1}{2}$ also have larger increases in their annual withdrawals in that year. Likewise, those with the least amount of exposure to penalty-free withdrawals in the $59\frac{1}{2}$ year experience the greatest increase in exposure when moving to the year in which they attain $60\frac{1}{2}$. Accordingly, those groups then experience the greatest growth in annual withdrawals. By the time age $61\frac{1}{2}$ is reached, the gap between the groups has decreased considerably, but not completely, suggesting that either higher withdrawals for those who turned $59\frac{1}{2}$ earlier persist to some extent after the penalty is lifted, or that the difference in age between these groups matters for withdrawals at these later ages. In general, we find a robust, monotonic relationship between penalty exposure and withdrawals.

As an alternative way to visualize this pattern, we show in Figure 1b average annual withdrawals for people grouped by their age during the final month of the year. We find a flat pattern of withdrawals to the left of $59\frac{1}{2}$, that is, for years where the early withdrawal penalty is in place for the entire year. For the group reaching age $59\frac{1}{2}$ during the final month of the year, the penalty is lifted for one month, and we see an upward shift in annual average withdrawals. This pattern continues, with incremental increases for each additional month of penalty-free withdrawals.

We generalize this notion more formally and develop a flexible, parametric approach for estimating the short-run response to the removal of the early withdrawal penalty. Consider people born on day b. Our goal is to estimate their average level of withdrawals (including zeros) at a daily frequency near the day on which age $59\frac{1}{2}$ is attained, which we denote as a_b^*—for example, $a_{b=30June1950}^* = 30Dec2009$. On a given "event day" j, i.e. j days relative to a_b^*, we model the daily withdrawals for people born on day b—denoted as $\tilde{y}_{bj}$—as follows:

$$\begin{aligned} \tilde{y}_{bj} = \tilde{\lambda}_{j+a_b^*} + & \underbrace{f_l(j) \cdot \mathbf{1}\{j < 0\} - D \cdot \mathbf{1}\{-w \leq j < 0\}}_{\text{left side of } a^*} \\ & + \underbrace{[L + f_r(j)] \cdot \mathbf{1}\{j \geq 0\} + S \cdot \mathbf{1}\{0 \leq j < w\}}_{\text{right side of } a^*} + \tilde{\varepsilon}_{bj} \end{aligned} \tag{1}$$

where $\tilde{\lambda}_d$ is a calendar-day fixed effect for day d, and $\tilde{\varepsilon}_{bj}$ is a mean-zero error term. On days earlier than age $59\frac{1}{2}$, that is on the "left" side of a^*, withdrawals are a smooth function of event time, $f_l(\cdot)$, with an additional, discrete dip, D, for dates within w days of a^*. Once the date a^* is reached, there is a permanent, longer-term shift in withdrawals, L, and we thereafter allow withdrawals to follow a different, smooth function of event time, $f_r(\cdot)$. We again allow for a discrete spike, S, within w days following a^*. The terms D and S allow for flexible, short-term re-timing of withdrawals.

To better match the structure of our tax data, we next aggregate the daily process in equation (1) to an annual frequency. We likewise define years in event time e, that is years relative to the year in which age $59\frac{1}{2}$ is attained. Let the mapping $t(b,e)$ be the calendar year in which someone born on day b reaches event year e—for example, $t(b = 30 June 1950, e = 0) = 2009$. Likewise, let $\mathcal{J}(b,e)$ be the set of event days j that occur during event year e for someone born on day b—for example, $\mathcal{J}(b = 30 June 1950, e = 0) = [-363, 1]$. Finally, for tractability, we model the smooth functions in event time as polynomials, i.e. $f(j) = \sum_p \beta_p j^p$. In Appendix E.1, we derive an expression for total annual withdrawals, by summing equation (1) over the event days in event year e:

$$
\begin{aligned}
y_{be} &\equiv \sum_{j \in \mathcal{J}(b,e)} \tilde{y}_{bj} \\
&= \lambda_{t(b,e)} + \sum_p \beta_{l,p} \left[\sum_{j \in \mathcal{J}(b,e)} j^p \cdot \mathbf{1}\{j < 0\} \right] - D \left[\sum_{j \in \mathcal{J}(b,e)} \mathbf{1}\{-w \leq j < 0\} \right] \\
&\quad + L \left[\sum_{j \in \mathcal{J}(b,e)} \mathbf{1}\{j \geq 0\} \right] + \sum_p \beta_{r,p} \left[\sum_{j \in \mathcal{J}(b,e)} j^p \cdot \mathbf{1}\{j \geq 0\} \right] + S \left[\sum_{j \in \mathcal{J}(b,e)} \mathbf{1}\{0 \leq j < w\} \right] + \varepsilon_{be}
\end{aligned}
\tag{2}
$$

Here $\lambda_{t(b,e)} \equiv \sum_{d \in t(b,e)} \tilde{\lambda}_d$ is a calendar-year fixed effect and $\varepsilon_{be} \equiv \sum_{j \in \mathcal{J}(b,e)} \tilde{\varepsilon}_{bj}$ is an annual, mean-zero error term. The resulting expression features the underlying polynomial terms and flexible shifters, $\{\beta_{l,1}, \ldots, \beta_{l,p}, \beta_{r,1}, \ldots, \beta_{r,p}, D, L, S\}$ as coefficients of terms that vary at the birthdate-by-event-year level, and can therefore be recovered via OLS regression using annual data.[5] The parameters are identified by variation in the set of event days j that occur within an event year e given one's birthdate b, including the number of days on either side of a^* or within w days of a^*, as previewed in Figure 1.

We estimate equation (2) using data collapsed at the date-of-birth-by-event-year cell, for a total of 14,608 observations. For ease of notation, we have shown uniform shifts (D, S) spanning the full window on either side of a^*. In practice, we divide each window into multiple, equally-sized steps.

Using the parameters estimated in equation (2), we can construct $\hat{y_{bj}}$, the predicted daily average withdrawal from equation (1).[6] We calculate the average daily withdrawal over an interval of event days, j_1 to j_2, as $\bar{y}_b(j_1, j_2) \equiv \sum_{j=j_1}^{j_2} \frac{\hat{y}_{bj}}{j_2 - j_1 + 1}$. We then define four measures of the short- and longer-run dynamics of withdrawal behavior surrounding age $59\frac{1}{2}$:

1. Longer-Run Change $\equiv \bar{y}_b(60, 89) - \bar{y}_b(-90, -61)$
2. Short-Run Change $\equiv \bar{y}_b(0, 29) - \bar{y}_b(-30, -1)$
3. Dip $\equiv \bar{y}_b(-90, -61) - \bar{y}_b(-30, -1)$
4. Spike $\equiv \bar{y}_b(0, 29) - \bar{y}_b(60, 89)$

The Longer-Run Change, which approximates L in equation (1), measures the change in average withdrawals between the third month prior to the penalty expiration and the third month following the penalty expiration. This will generally abstract from shorter-run dynamics and retiming that are reflected in our Short-Run Change measure, which captures the average change in withdrawals from the month before to the month after the early withdrawal penalty expires. This short-run retiming has two components. First, the "Dip", or D in equation (1), captures any decline in withdrawals as the penalty expiration nears, presumably reflecting a delay in withdrawals until the penalty is gone. The "Spike",

[5]Note that our proposed methodology is unable to recover the *share* of people taking a withdrawal on a given day, as the annual data cannot distinguish between someone who takes one withdrawal and someone who takes many withdrawals in a given year. Instead, the estimate would recover the share of people on a given day making their *first* withdrawal that year. This daily share aggregates at the annual level to measure the share of people who made at least one withdrawal during the year.

[6]This predicted mean is *net* of the daily fixed effects and error term, i.e. $\hat{y}_{bj} = y_{bj} - \tilde{\lambda}_{j+a_b^*} - \tilde{\varepsilon}_{bj}$

or S in equation (1), reflects excess withdrawals that occur immediately after the penalty is removed, capturing any delayed withdrawals from the Dip region, but also potentially including additional very short-run increases in withdrawals. Note, we have the following relationship between the measures: Longer-Run Change = Short-Run Change - (Dip + Spike).

We estimate the average level of withdrawals at baseline as $\bar{y}_b(-90, -61)$. We also construct the share of baseline withdrawals that are penalized, which we denote θ, by estimating equation (2) using penalties paid on early withdrawals as the outcome variable. We then determine the baseline average level of penalties over the 60-90-day interval prior to the penalty removal, dividing it by the penalty rate of 10 percent to get the amount of withdrawals subject to a penalty, and dividing that amount by the baseline level of withdrawals.

The above discussion asserts that we can recover a parametric approximation to the daily pattern of average withdrawals with only annual data and variation in day of birth across tax filers. We use simulations to validate our method. Appendix B features figures showing simulated daily patterns of withdrawals, which are then aggregated to annual amounts. We then show the daily pattern that is recovered using our estimator on the aggregated, simulated data. Returning to Figure 1b, we see that the increase in average annual withdrawals after the first month of no penalties is more than half the height of the increase after the first two months of no penalties. That ratio of the former to the later is 0.59. Intuitively, these moments help us identify the additional increase in withdrawals during the first post-penalty month, relative to second post-penalty month; that is, we find patterns related to the presence of the short-term "spike."[7]

Provided that we can, in fact, recover the daily patterns in withdrawals, there is still a question of interpretation. In order to ascribe changes in behavior near age $59\frac{1}{2}$ to the removal of the early withdrawal penalty, we must make assumptions about the counterfactual pattern of withdrawals in the absence of the policy change. As is common in event study

[7] In simulations, available from the authors upon request, we confirm that the ratio of these two moments is indeed increasing in the ratio of the spike to the longer-run increase in withdrawals, all things equal.

designs, we assume that if the penalty were to remain at age $59\frac{1}{2}$, the patterns of withdrawals would continue to evolve locally according to the smooth function $f_l(j)$.

4 Results

4.1 Main results

We trace out daily IRA withdrawal patterns before and after the 10 percent penalty for early withdrawals is lifted at age $59\frac{1}{2}$ using the approach outlined above. We present estimates of the withdrawal patterns graphically in Figure 2 using a third degree polynomial and a non-parametric window of 60 days on either side of $59\frac{1}{2}$, with each window divided into three equally-sized steps. We report the longer-run effect, short-run effect, dip, and spike, along with the baseline withdrawals and the share of baseline withdrawals that are penalized, in the first row of Table 1.[8]

Our estimates suggest that the increase in daily withdrawals over the longer-run is \$5.14 per day or \$1,876 per year. This increase is large compared to the baseline level of withdrawals of \$4.93 per day or \$1,799 per year, amounting to a doubling of withdrawals after the penalty is lifted. Our empirical approach allows us to understand whether this longer-term increase masks any short-term dynamics in the days just before and after $59\frac{1}{2}$.

As shown in Figure 2, the evidence suggests that the jump in withdrawals occurs rapidly after individuals turn $59\frac{1}{2}$, with the level stabilizing approximately 30 days afterwards. We do not find strong evidence that withdrawals decline just prior to age $59\frac{1}{2}$. However, our results do indicate that withdrawals spike immediately after age $59\frac{1}{2}$, before decreasing to a level still higher than at baseline. As shown in Table 1, the estimated dip in withdrawals is not statistically significant, while the estimated spike in withdrawals in the 30-day period following the penalty's removal is \$6.62 per day, above and beyond the long-term increase

[8] As shown in Appendix D, these results are not sensitive to the choice of parameters, so we focus on one set of results for expository purposes.

of \$5.14 per day. This temporary spike accounts for more than half of the short-run change in withdrawals.

We investigate the robustness of our results in two ways. First, we perform a set of placebo tests using alternative age thresholds, limited to data that do not cross the $59\frac{1}{2}$ threshold and thus are not contaminated by any true increase in withdrawals that may be happening at that age. We estimate short-run changes in withdrawals that are clustered around zero and summarize these results along with our estimated effect at age $59\frac{1}{2}$ in Figure 3a. Second, we perform a permutation exercise where we use the full set of data and cycle through age thresholds ranging from age 58 and zero months to 61 and zero months. We display these results in Figure 3b and find that the short-run effect estimated using these different age thresholds are quantitatively small and generally insignificant. Together, these results reinforce the notion that the effects we estimate at age $59\frac{1}{2}$ are not a mechanical result of our empirical method, but rather the result of removing the penalty for early withdrawals at age $59\frac{1}{2}$.

4.2 Heterogeneity

We estimate heterogeneity in the short-term and longer-term effects of the penalty's removal across observable characteristics and time periods to understand how our results vary when liquidity constraints are more likely to be binding, and whether the responses we estimate are present at different parts of the income and wealth distribution. We split our sample by quartile of AGI and FMV of each person's IRA accounts in the year in which they have reached $57\frac{1}{2}$. We also examine those who turned $59\frac{1}{2}$ during the Great Recession between December 2007 and June 2009, those who were receiving unemployment insurance (UI) at age $57\frac{1}{2}$, and those who were receiving income from the Social Security Disability Insurance (SSDI) program at age $57\frac{1}{2}$. These latter three groups comprise 18 percent, 5 percent, and 3 percent of the observations, respectively. The estimates of these analyses are reported for each of these groups in rows 2-6 of Table 1.[9]

The average gross income in the top quartile of AGI in our sample at age $59\frac{1}{2}$ is $307,119 owing to the relatively high income in our IRA-holding sample. This group experiences higher levels of baseline withdrawals, and proportionally higher short-run increases at age $59\frac{1}{2}$, compared to the full sample. Those with FMVs in the highest quartiles have IRA balances of $304,296 on average, and an average AGI of $171,473. This group has baseline withdrawals about three times as large as the full sample, and proportionally similar increases in the short run. These results suggest that the short-run response we detect is not only present in lower segments of the income/wealth distribution, but also occurs among relatively high-income and wealthy populations.

For the subsamples turning $59\frac{1}{2}$ during the Great Recession or with a recent history of receiving UI, we make two important observations. First, a high share of baseline withdrawals are penalized, suggesting that liquidity constraints prior to age $59\frac{1}{2}$ are leading individuals to make early withdrawals despite the 10 percent penalty. Second, the proportional increases in the short run are 4-5 times higher than baseline withdrawals for these groups, compared to 2.36 times higher for the full sample.

Finally, we examine those who recently received SSDI benefits in the last column of Table 1. Despite recently experiencing a negative income shock, the results for this group differ from the results for the Great Recession and UI subsamples. While we estimate a longer-run increase in withdrawals that is similar in magnitude to the full sample, the short-run increase and spike are both statistically insignificant. Even with a large standard error, the upper bound on the 95 percent confidence interval for the SSDI group's estimated spike is roughly half of the analogous coefficient estimated in the UI and Great Recession samples.

These results are consistent with the IRS's broad penalty exception for withdrawals made on account of disability. Indeed, we estimate that a lower share of baseline withdrawals were subject to the penalty prior to age $59\frac{1}{2}$ for this group. While exceptions are also

[9] Appendix C features results for all four quartiles of AGI and FMV.

available during unemployment, the penalty is only waived for the unemployed to cover health insurance premiums. The results in the last two columns therefore highlight the differences in withdrawal patterns between subsamples of individuals experiencing significant shocks to income, one which has limited access to IRA account balances and one for whom IRA balances provide an important source of liquidity following a negative income shock.

5 Discussion

We find evidence of elevated withdrawals from IRAs immediately following the removal of the early withdrawal penalty at age $59\frac{1}{2}$. Approximately 56 percent of the higher level of withdrawals we estimate within 30 days of the penalty expiration represents a temporary spike in withdrawals, while 44 percent persists over the longer term, beyond this 30-day window. In this section, we demonstrate how these empirical estimates might serve as inputs into an assessment of the welfare impact of two policy proposals: (1) an adjustment in the age at which the early withdrawal penalty applies, e.g. a shift of the threshold to x days earlier, and (2) a temporary penalty holiday, allowing penalty-free withdrawals to be made for a period of x days. One approach to evaluating the policy is to calculate the marginal value of public funds (MVFP), as in Hendren and Sprung-Keyser (2020). This entails dividing the willingness to pay for the policy change by its net fiscal cost.

First, suppose the age at which withdrawals no longer face penalties is shifted from the day when age $59\frac{1}{2}$ is attained, a^*, to x days earlier, or day $a^* - x$. We begin by describing the net fiscal cost in the denominator of the MVPF. On each day between $a^* - x$ and a^*, a daily average of $With_0$ in withdrawals would have been made even in the absence of the policy change, which we refer to as inframarginal withdrawals. For each inframarginal dollar, a penalty of $\pi = 10\%$ is forgone, from the government's perspective. Assuming θ share of dollars would have been subject to the penalty, the resulting cost is then $\pi\theta \times With_0$ per day, for x days.

In addition, there is a marginal increase in average withdrawals, $\triangle With$, which we estimate as the behavioral response to the penalty removal at age $59\frac{1}{2}$. In our setting, this increase entails a longer-run increase in withdrawals, or the parameter L in equation (1), plus an additional "spike" in average daily withdrawals during the first month after age $59\frac{1}{2}$, or the parameter S in equation (1).[10] For ease of argument, we assume $x > 30$, so that the "spike" lasts for the first 30 days, and the longer-run increase lasts during the entire time interval.

We assume that in the absence of this policy change, these marginal withdrawals would have happened, as we observe, x days later, after age $59\frac{1}{2}$. Since no penalties are paid in either case on these marginal withdrawals, there is no change in government revenue from penalties. In addition, because the funds are taken out of a tax-preferred account earlier, they represent a reduction in government-subsidized returns, and thus a marginal increase in tax revenue. In Appendix E.2, we derive an expression for this foregone subsidy: $\tau_r r (1 - \tau) \times \triangle With$. In short, this is the income tax of τ_r, on a return of r, that would have been earned during the additional x days that these funds would have been held in an IRA, applied to an after-income-tax principal of $(1 - \tau) \times \triangle With$. During the first 30 days, $\triangle With = L + S$, and during the remainder, $\triangle With = L$. This amount is discounted to the present by a factor of δ.

In this baseline case, the numerator of the MVPF equals the value of avoided penalties on inframarginal withdrawals, or $\pi\theta \times With_0$ per day. This assumes that the value of the additional marginal withdrawals, in response to earlier penalty expiration, is roughly equal to the discounted value of forgone withdrawals that would have happened upon reaching age $59\frac{1}{2}$, an application of the envelope theorem. Combining the above components, we arrive

[10] There is potentially an additional short-run retiming, as captured by a "dip" in withdrawals just prior to the expiration of the penalty. However, since we do not find robust evidence of this response, we ignore it for the sake of this exercise.

at the following MVPF of a shift of x days in the age at which the penalty expires:

$$
\begin{aligned}
MVPF_{\text{SHIFT}} &= \frac{\pi\theta With_0 \times x}{\pi\theta With_0 \times x - \delta\tau_r r\,(1-\tau)\,[L \times x + S \times 30]} \\
&= \frac{\pi\theta}{\pi\theta - \delta\tau_r r\,(1-\tau)\,[\alpha_L + \alpha_S]}
\end{aligned} \tag{3}
$$

where $\alpha_L \equiv \frac{L \times x}{With_0 \times x}$ is the proportional increase in withdrawals attributable to the longer-term response, and $\alpha_S \equiv \frac{S \times 30}{With_0 \times x}$ is that attributable to the short-term spike. In this baseline case, the MVPF is greater than one: the foregone penalties are a transfer from the government to the savers, and the additional behavioral responses result in fewer tax expenditures on tax-preferred savings.

In Table 1, we estimate values of the key parameters: $\theta \approx 0.23$, $\alpha_L \approx 1.04$, and $\alpha_S \approx 1.34 \times \frac{30}{x}$. If we shift the penalty expiration a year earlier, we have $30/x \approx 1/12$, and, assuming an annual, before-tax, return on savings of $r = 0.06$, marginal tax rates on capital and labor income of $\tau_r = \tau = 0.3$, an annual discount factor of $\delta = 0.97$, and input the early withdrawal penalty of $\pi = 0.1$, we have $MVPF_{\text{Shift}} = 2.58$. This implies that every \$1 spent by the government on moving the penalty age earlier delivers \$2.58 in benefits to those who withdraw early. A shift of the penalty expiration to an older age would symmetrically yield $-MVPF_{\text{Shift}}$.

Next, we consider a penalty holiday that lasts for x days, at an arbitrary age, potentially many years younger than age $59\frac{1}{2}$. We interpret the longer-term change in withdrawals after age $59\frac{1}{2}$ as a response to the permanent removal of the penalty, and therefore assume that there is only the short-term response to this temporary removal, i.e. the spike in withdrawals, S, during the first 30 days of this holiday. Relative to the expression in equation (3), we therefore remove the α_L term in the denominator:

$$
MVPF_{\text{HOLIDAY}} = \frac{\pi\theta}{\pi\theta - \delta\tau_r r\,(1-\tau)\,\alpha_S} \tag{4}
$$

We again assume that the marginal withdrawal would have occurred at age $59\frac{1}{2}$ in the absence of this holiday. In addition, given the greater amount of time between the holiday and age $59\frac{1}{2}$, the discount factor δ will tend to be smaller, and the value of tax subsidy, $\tau_r r (1-\tau)$ will tend to be larger, since the returns are allowed to compound over a longer period.[11] For example, a one-year holiday occurring 5 or 10 years earlier than age $59\frac{1}{2}$ results in an MVPF of 1.47 or 3.41, respectively. Supposing instead that the holiday occurs 20 years earlier, the value of the deferred tax payments dominates, the denominator becomes negative, and, following Hendren and Sprung-Keyser (2020), we set $MVPF_{\text{Holiday}}$ to equal ∞. In other words, the savings from a decrease in deferred earnings more than pays for the loss of early withdrawal penalties.

The above exercises abstract from a number of factors. Importantly, the net impact of marginal early withdrawals on utility disappears in the numerator of the MVPF, due to the envelope theorem. However, if there are wedges between the private and social payoff of induced retirement savings, this additional factor would be included in the numerator. The original presence of a penalty, and tax-preferences for retirement savings may be indicative of such a wedge; however, the implications for the MVPF depend on the nature of the wedge. For example, if there are externalities imposed on society from undersaving, a negative term may be added to the numerator, leading to a smaller MVPF. Moreover, behavioral models that incorporate time inconsistency may increase or decrease the numerator of the MVPF, depending on whether the present or future self's preferences are valued more, respectively. This is similarly the case if some of the additional withdrawals occur due to a relaxation of binding borrowing constraints, in which case the envelope theorem may not apply. This case is also likely to include a more positive numerator and larger MVPF. Finally, if early withdrawals are spent on goods and services, this may generate a positive multiplier effect on the economy, particularly during an economic downturn. Such a multiplier effect would further increase the MVPF, making these policies potentially more beneficial in an economic

[11] See Appendix E.2, for derivations of exact expressions. With returns r_t and capital income taxes τ_t in each period t: $r \equiv \prod_{t=1}^{t^*}(1+r_t)-1$ and $\tau_r \equiv \left[\prod_{t=1}^{t^*}(1+r_t)-\prod_{t=1}^{t^*}(1+r_t(1-\tau_t))\right] \Big/ \left[\prod_{t=1}^{t^*}(1+r_t)-1\right]$.

downturn.

Our empirical estimates only apply locally to those near age $59\frac{1}{2}$. As such, applying our estimates to the behavior of those significantly younger than age $59\frac{1}{2}$, as in the case of a penalty holiday, involves nontrivial extrapolation. This is less of a concern in the case of a local shift of the age at which the penalty expires. Finally, if either of these policies were to become permanent, there may be corresponding behavioral changes in deposits to retirement savings accounts that we do not capture, but that would matter for the calculation of the MVPF.

6 Conclusion

We examine how withdrawals from IRAs change when the penalty for early withdrawals is lifted upon the account holder turning age $59\frac{1}{2}$. We develop a method to use annual data to uncover patterns at a subannual frequency and find robust evidence that withdrawals increase substantially after penalties are removed, and that the short-run increase in withdrawals is significantly larger than the long-run increase. On average, withdrawals increase by 3 to 3.5 times the baseline amount within one month of the penalty removal, and decline in the longer-run to a doubling of the baseline amount within three months. These findings are present across all quartiles of income and account balances, and are more pronounced for recently unemployed individuals and those turning $59\frac{1}{2}$ during the Great Recession, when liquidity constraints are more likely to be binding. However, those who have received disability income do not exhibit a large short-run increase, likely because they qualify for an exception to penalties on early withdrawals prior to age $59\frac{1}{2}$.

We further describe how our estimates might be used to construct estimates of the marginal value of public funds for different policy proposals that affect the liquidity of tax-preferred retirement savings accounts. Our framework highlights the importance of understanding the relative magnitudes of the short-term and longer-term effects of increased liquidity in these accounts. While combining our empirical findings with such a framework suggests that such policies as an early age of penalty expiration or a temporary penalty holiday may have positive welfare impacts, an alternative set of assumptions may yield different conclusions.

References

Agarwal, Sumit, Jessica Pan, and Wenlan Qian. 2020. "Age of decision: Pension savings withdrawal and consumption and debt response." *Management Science*, 66(1): 43–69.

Amromin, Gene, and Paul Smith. 2003. "What Explains Early Withdrawals From Retirement Accounts? Evidence From a Panel of Taxpayers." *National Tax Journal*, 595–612.

Argento, Robert, Victoria L Bryant, and John Sabelhaus. 2015. "Early Withdrawals From Retirement Accounts During the Great Recession." *Contemporary Economic Policy*, 33(1): 1–16.

Beshears, John, James J Choi, Joshua Hurwitz, David Laibson, and Brigitte C Madrian. 2015. "Liquidity in Retirement Savings Systems: An International Comparison." *The American Economic Review*, 105(5): 420–425.

Brown, Jeffrey R, James Poterba, and David Richardson. 2017. "Do Required Minimum Distributions Matter? The Effect of the 2009 Holiday On Retirement Plan Distributions." *Journal of Public Economics*, 151: 96–109.

Coyne, David, Itzik Fadlon, and Tommaso Porzio. 2020. "Who Needs Liquidity, When, and Where? Evidence from Penalized Withdrawals." *Working Paper.*

Deryugina, Tatyana, and Solomon M Hsiang. 2014. "Does the Environment Still Matter? Daily Temperature and Income in the United States." *NBER Working Paper 20750.*

Deryugina, Tatyana, Laura Kawano, and Steven Levitt. 2018. "The Economic Impact of Hurricane Katrina on Its Victims: Evidence from Individual Tax Returns." *American Economic Journal: Applied Economics*, 10(2): 202–33.

Deschênes, Olivier, and Michael Greenstone. 2011. "Climate Change, Mortality, and Adaptation: Evidence From Annual Fluctuations in Weather in the US." *American Economic Journal: Applied Economics*, 3(4): 152–185.

Goodman, Lucas, Kathleen Mackie, Jacob Mortenson, and Heidi Schramm. 2019. "The Evolution of Leakage and Retirement Asset Flows in the US." *Available at SSRN 3450866.*

Hendren, Nathaniel, and Ben Sprung-Keyser. 2020. "A unified welfare analysis of government policies." *The Quarterly Journal of Economics*, 135(3): 1209–1318.

Horneff, Vanya, Raimond Maurer, and Olivia S. Mitchell. 2021. "Do Required Minimum Distribution 401(k) Rules Matter, and For Whom? Insights from a Lifecycle Model." National Bureau of Economic Research Working Paper 28490.

Mortenson, Jacob A., Heidi R. Schramm, and Andrew Whitten. 2019. "The Effect of Required Minimum Distribution Rules on Withdrawals from Traditional Individual Retirement Accounts." *National Tax Journal*, 72(3): 507–542.

Poterba, James, Steven Venti, and David Wise. 2013. "The Drawdown of Personal Retirement Assets: Husbanding or Squandering?"

Stuart, Ellen, and Victoria Bryant. 2020. "The Impact of Withdrawal Penalties on Retirement Savings." *Working Paper*.

Exhibits

Figure 1: Average Annual IRA Withdrawals by Exposure to Penalty-Free Withdrawal

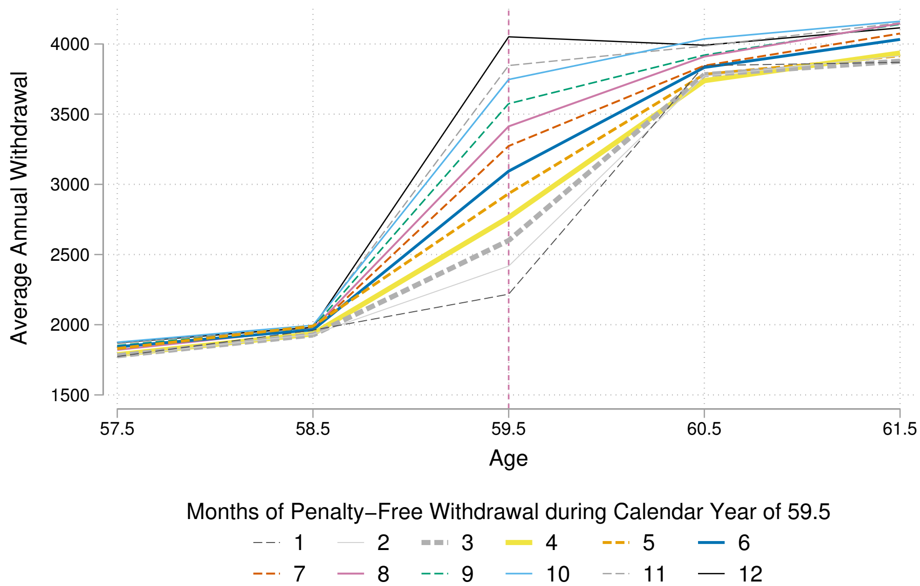

(a) Withdrawals over time

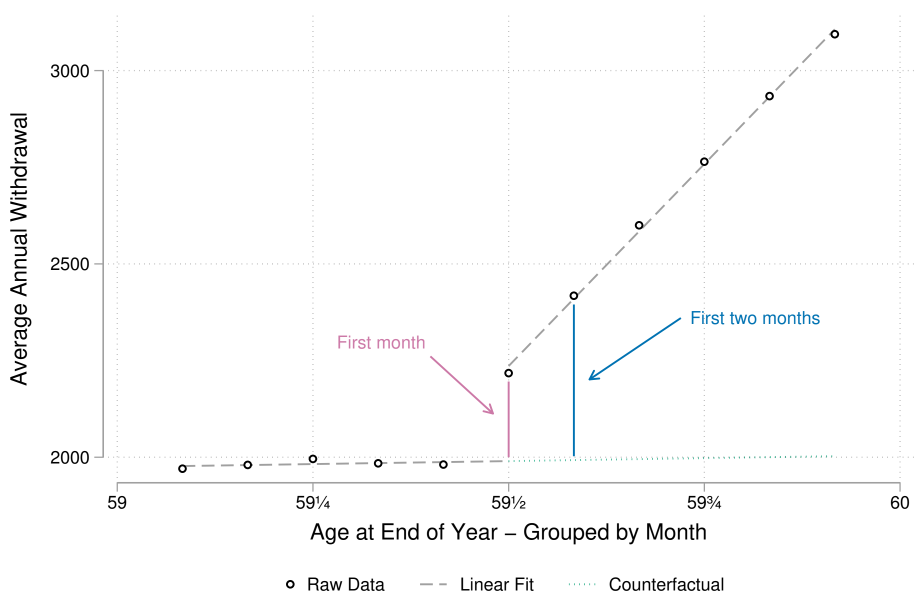

(b) Withdrawals by age at end of year

Notes: Panel (a) shows average annual IRA withdrawals across birth cohorts born in different calendar months during the calendar years they attain age $57\frac{1}{2}$, $58\frac{1}{2}$, $59\frac{1}{2}$, $60\frac{1}{2}$ and $61\frac{1}{2}$. Panel (b) shows average annual withdrawals for people grouped by their age in December of that year.

Figure 2: Estimated Average Daily IRA Withdrawals

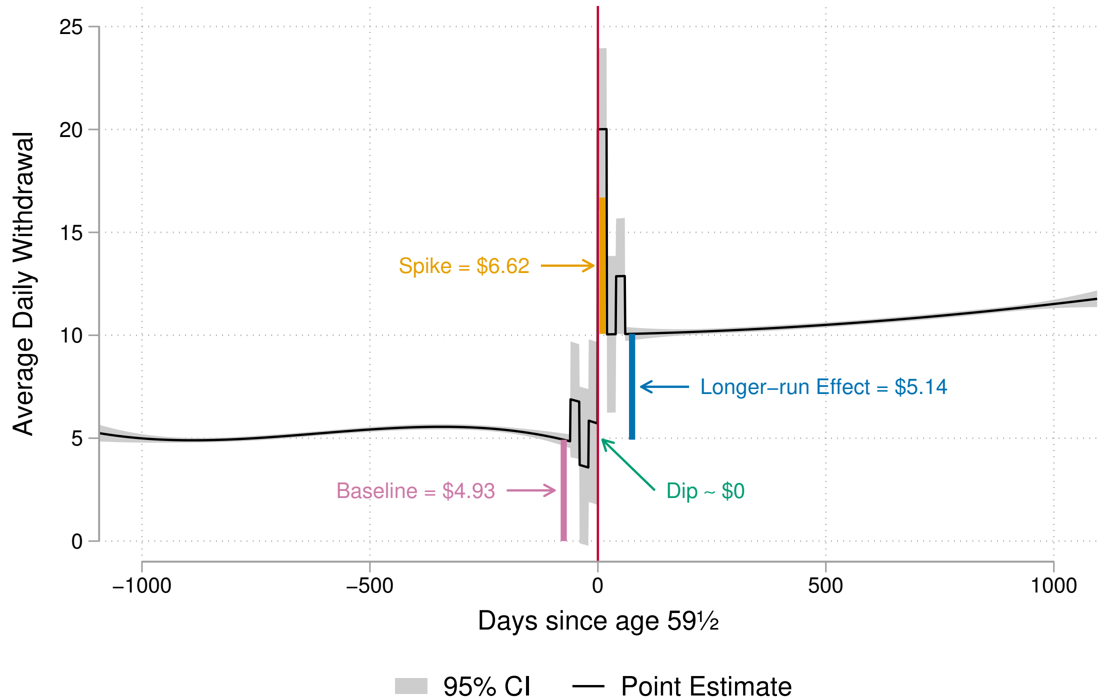

Notes: Figure shows predicted values of average daily IRA withdrawals from estimating equation 2 with polynomial order 3, 60 day window, and 3 non-parametric steps.

Figure 3: Results of Placebo and Permutation Tests

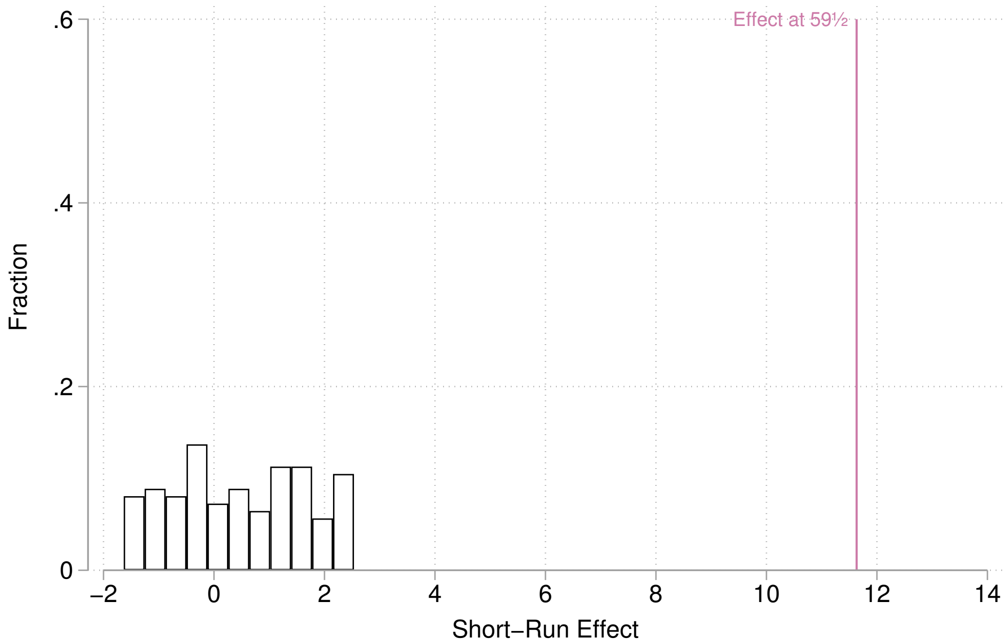

(a) Placebo estimates of short-run effect

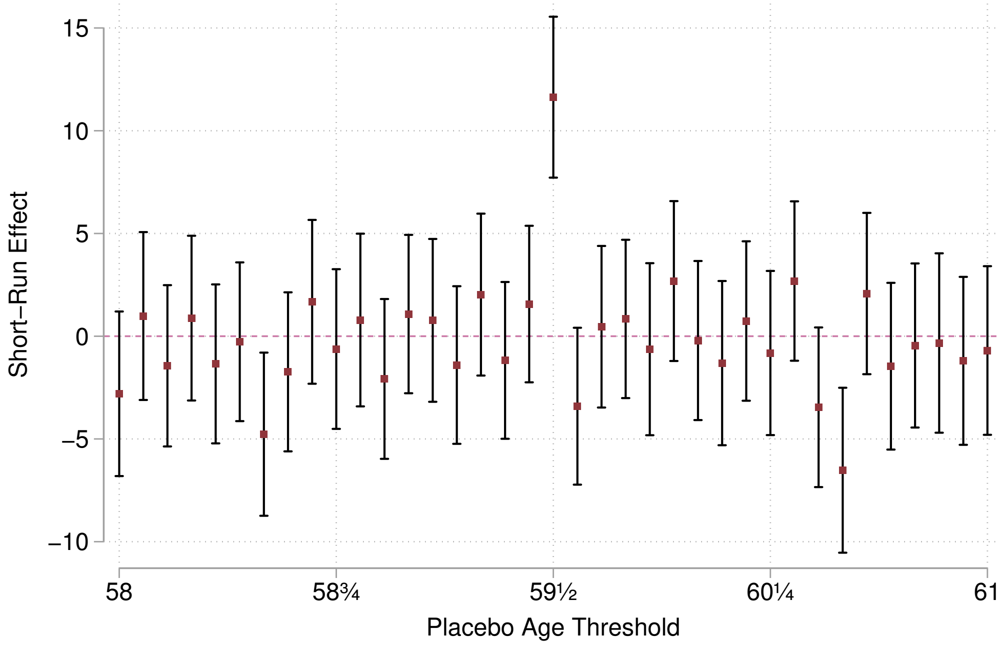

(b) Short-run effect measured for alternative age thresholds

Notes: Top panel shows histogram of placebo values of Short-Run Effect across placebo age thresholds [57 and 9 months, 57 and 11 months] using data to the left of the age $59\frac{1}{2}$ threshold and [60 and 11 months, 61 and 1 month] using data to the right of the age $59\frac{1}{2}$ threshold. Bottom panel shows estimated Short-Run Effect using full data and alternative age thresholds spanning [58 and 0 months, 61 and 0 months].

Table 1: Changes in IRA Daily Withdrawals at Age $59\frac{1}{2}$

	(1) Full Sample	(2) Top 75^{th} AGI	(3) Top 75^{th} FMV	(4) Great Recession	(5) Claims UI	(6) Claims DI
Short-Run	11.63	20.36	30.08	23.94	17.89	-2.24
	(2.00)	(6.30)	(8.02)	(5.63)	(7.18)	(8.38)
Longer-Run	5.14	5.39	12.41	5.72	8.61	4.63
	(0.19)	(0.59)	(0.75)	(0.55)	(0.67)	(0.82)
Dip	-0.13	6.09	-2.79	4.46	-4.01	-5.68
	(1.07)	(3.36)	(4.28)	(3.16)	(3.83)	(4.41)
Spike	6.62	8.88	20.45	13.76	13.29	-1.19
	(1.06)	(3.35)	(4.26)	(2.89)	(3.82)	(4.56)
Baseline Withdrawals	4.93	7.64	15.60	4.21	4.47	5.54
Share Penalized	0.229	0.289	0.152	0.384	0.379	0.104
Person Obs	12,445,149	3,111,234	3,111,218	2,937,441	579,539	425,332

SEs are in parentheses.

Notes: Standard errors calculated using the delta method. Short-Run and Long-Run Effects represent increase in average IRA daily withdrawals over 1 month before and after age $59\frac{1}{2}$, and from 2-3 months before age $59\frac{1}{2}$ to 2-3 months after age $59\frac{1}{2}$, respectively. Dip represents reduction in average IRA daily withdrawals from 2-3 months before age $59\frac{1}{2}$ to 0-1 month before age $59\frac{1}{2}$, and Spike represents increase in average IRA daily withdrawals from 2-3 months after age $59\frac{1}{2}$ to 0-1 month after age $59\frac{1}{2}$. Baseline represents average level of daily IRA withdrawals 2-3 months prior to age $59\frac{1}{2}$.

Appendix A: Data Appendix

To construct our primary dataset, we start with the Social Security Administration's Death Master File (DM1), which comprises every individual in the United States who applies for either a Social Security Number (SSN) or an individual taxpayer identification number (ITIN), and provides exact dates of birth and dates of death. We then merge to that by SSN, information from tax and information returns. We restrict our attention to all individuals born between July 1, 1941 and July 1, 1951 who are still alive at age $57\frac{1}{2}$. In total, there are 34,423,900 people who meet those criterion. We further restrict that sample to people who have an individual retirement account in their own name with a positive fair market value in their year they turned $57\frac{1}{2}$. This restriction leaves 12,445,149 individuals for whom we pull the full set of tax variables used in our analysis. Finally, we drop 168 individuals from our estimation due to errors in the reporting of their date of birth.

The IRS calculates age $59\frac{1}{2}$ by determining the month and year in which an individual turns 59, moving six months forward, and then choosing the day in that month that corresponds to the day of birth.[12]

Our data on IRA distributions comes from Form 1099-R, which reports Distributions From Pensions, Annuities, Retirement or Profit-Sharing Plans, IRAs, Insurance Contracts, etc. We use the "IRA/SEP/SIMPLE" checkbox on the form to identify distributions from IRAs accounts and further restrict our attention to early and normal distributions from Traditional IRAs using Box 7 codes: 1,2,7,D, L. These codes allow us to exclude distributions due to death, disability, rollovers, or conversions that are unlikely related to the removal of the penalty. We additionally exclude Roth IRA distributions, which are made with after-tax income and therefore, subject to different rules pertaining to the penalty. Data for the penalty variable comes from Form 1040 and gives the additional tax on IRAs, other qualified

[12]While in most cases this is straightforward, there are some cases where special rules apply. For instance, if someone is born on August 31, 1970 the above rules would specify February 31, 2030 as the day they turn $59\frac{1}{2}$. Since this day does not exist, the rules specify to calculate the residual days left over at the end of the month and advance that many days forward (March 3, 2030 in this example). Individuals born on leap days turn $59\frac{1}{2}$ on September 1 in the year in which they turn 59.

plans, etc.

We also draw a sample of respondents from the waves 3-11 (1996-2012) of the Health and Retirement Study (HRS) that mimics our sampling frame in the IRS data by selecting respondents who have an IRA account when they are within one year of turning $57\frac{1}{2}$. Because the HRS asks about distributions made since the last interview, we randomly assign a withdrawal date over the reference period and aggregate these randomly assigned withdrawals over a calendar year to construct summary statistics that are comparable to those from the IRS data.

Table A.1: IRS Summary Statistics

	(1) 57.5	(2) 58.5	(3) 59.5	(4) 60.5	(5) 61.5
Fraction Male	0.49	0.49	0.49	0.49	0.49
Fraction Married	0.74	0.73	0.73	0.72	0.71
Adjusted Gross Income	134841.80	128522.28	124190.44	120452.14	117819.01
IRA Distribution	1827.22	1967.41	3188.56	3903.07	4081.25
IRA Distribution (non-zero)	25263.76	23552.38	23555.08	23933.24	23898.01
Fraction with IRA Distribution	0.07	0.08	0.13	0.16	0.17
IRA Fair Market Value	96802.69	101035.20	108727.61	112359.30	120447.29
Receive SSDI/SSI $57\frac{1}{2}$	0.03	0.03	0.03	0.03	0.03
Receive UI at $57\frac{1}{2}$	0.05	0.05	0.05	0.05	0.05
Observations	12445149	12445149	12445149	12445149	12445149

Notes: This table provides summary statistics for the population of individuals born between July 1, 1941 and July 1, 1951, who have own a non-zero balance Individual Retirement Account in the year they turn $57\frac{1}{2}$. Values are inflation-adjusted using the Consumer Price Index to 2010 dollars. Data Source: U.S. tax and information records.

Table A.2: HRS Summary Statistics

	(1) 57.5	(2) 58.5	(3) 59.5	(4) 60.5	(5) 61.5
Fraction Male	0.47	0.48	0.48	0.50	0.48
	(0.50)	(0.50)	(0.50)	(0.50)	(0.50)
Fraction Married	0.78	0.78	0.80	0.79	0.81
	(0.41)	(0.42)	(0.40)	(0.41)	(0.39)
HH income	166453	144174	109869	97868	102222
	(617248)	(475634)	(187597)	(109237)	(125743)
IRA Distribution (non-zero)	20000	19066	23937	24604	26899
	(47790)	(31449)	(42425)	(41995)	(45869)
Fraction with IRA Distribution	0.11	0.10	0.14	0.14	0.09
	(0.32)	(0.30)	(0.35)	(0.35)	(0.28)
IRA Account Balance	164013	168997	179338	169135	170075
	(298212)	(316939)	(480038)	(292323)	(264863)
Receive SSDI	0.03	0.04	0.06	0.02	0.03
	(0.16)	(0.20)	(0.25)	(0.15)	(0.18)
Receive UI	0.04	0.04	0.03	0.03	0.03
	(0.21)	(0.19)	(0.18)	(0.17)	(0.17)
Observations	2177	2185	2358	2250	2331

Notes: HRS summary statistics reported for sample of IRA owners within one year of turning $57\frac{1}{2}$ during waves 3-11 (1996-2012) and are constructed using weights. IRA withdrawals are randomly assigned to occur over the reference period. Values are inflation-adjusted using the Consumer Price Index to 2016 dollars. Prior to 2000, SSDI receipt indicator may also include SSI recipients.

Appendix B: Simulated data

We use a simulation to validate our method. First, we simulate patterns of IRA withdrawals at the daily level using the data generating process in (1) and aggregate the daily data at the calendar-year level to mimic the information provided from annual IRS data. We then use our method to recover the parameters of the original simulation. We explore the robustness of our method by using a range of parameters, varying the polynomial order, the size of the non-parametric window, and the number of non-parametric steps dividing the window. We simulate both a model with no response to the penalty expiration (Sim 1) and one with dynamics similar to what we find in the true data (Sim 2). The estimator performs well, finding no effects in our null simulation and picking up both long- and short-run changes in withdrawal patterns when they are in fact included in the underlying simulation (see Figure B.1). The measures defined above (Long-Run, Short-Run, Dip, Spike) are stable across a wide range of parameters provided the number of steps within a window is three or more (see Figure B.2).

Figure B.1: Simulated Data: True Daily Patterns and Estimated Event Studies

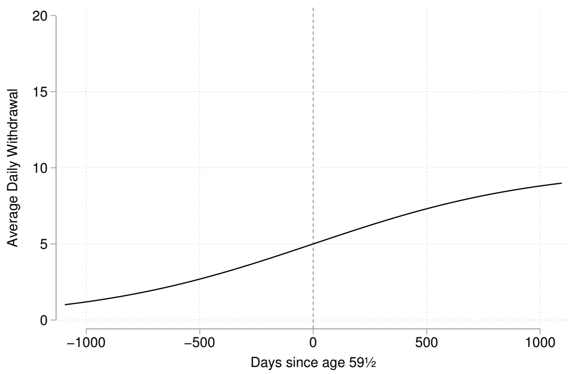

(a) Daily data: Sim 1

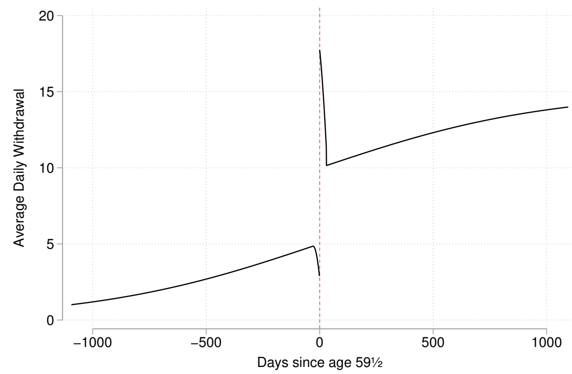

(b) Daily data: Sim 2

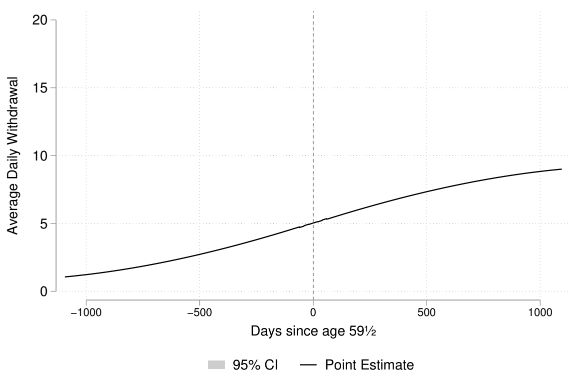

(c) Estimated event study: Sim 1

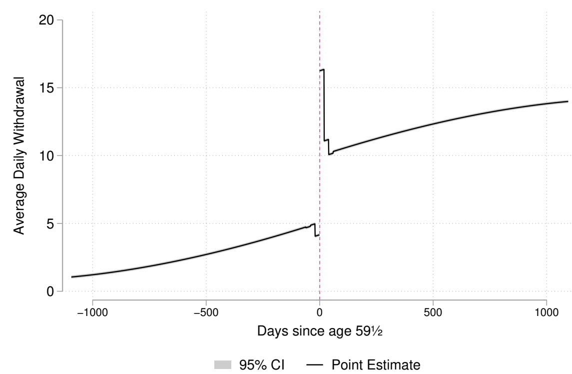

(d) Estimated event study: Sim 2

Notes: Estimated event studies use 3 steps, 60-day window, and 3rd degree polynomial. Sim 1 true short-run effect = 0.15; estimated short-run effect = 0.15 (0.00). Sim 2 true short-run effect = 10.72; estimated short-run effect = 10.18 (0.00).

Figure B.2: Simulated Data: Sensitivity to Parameters

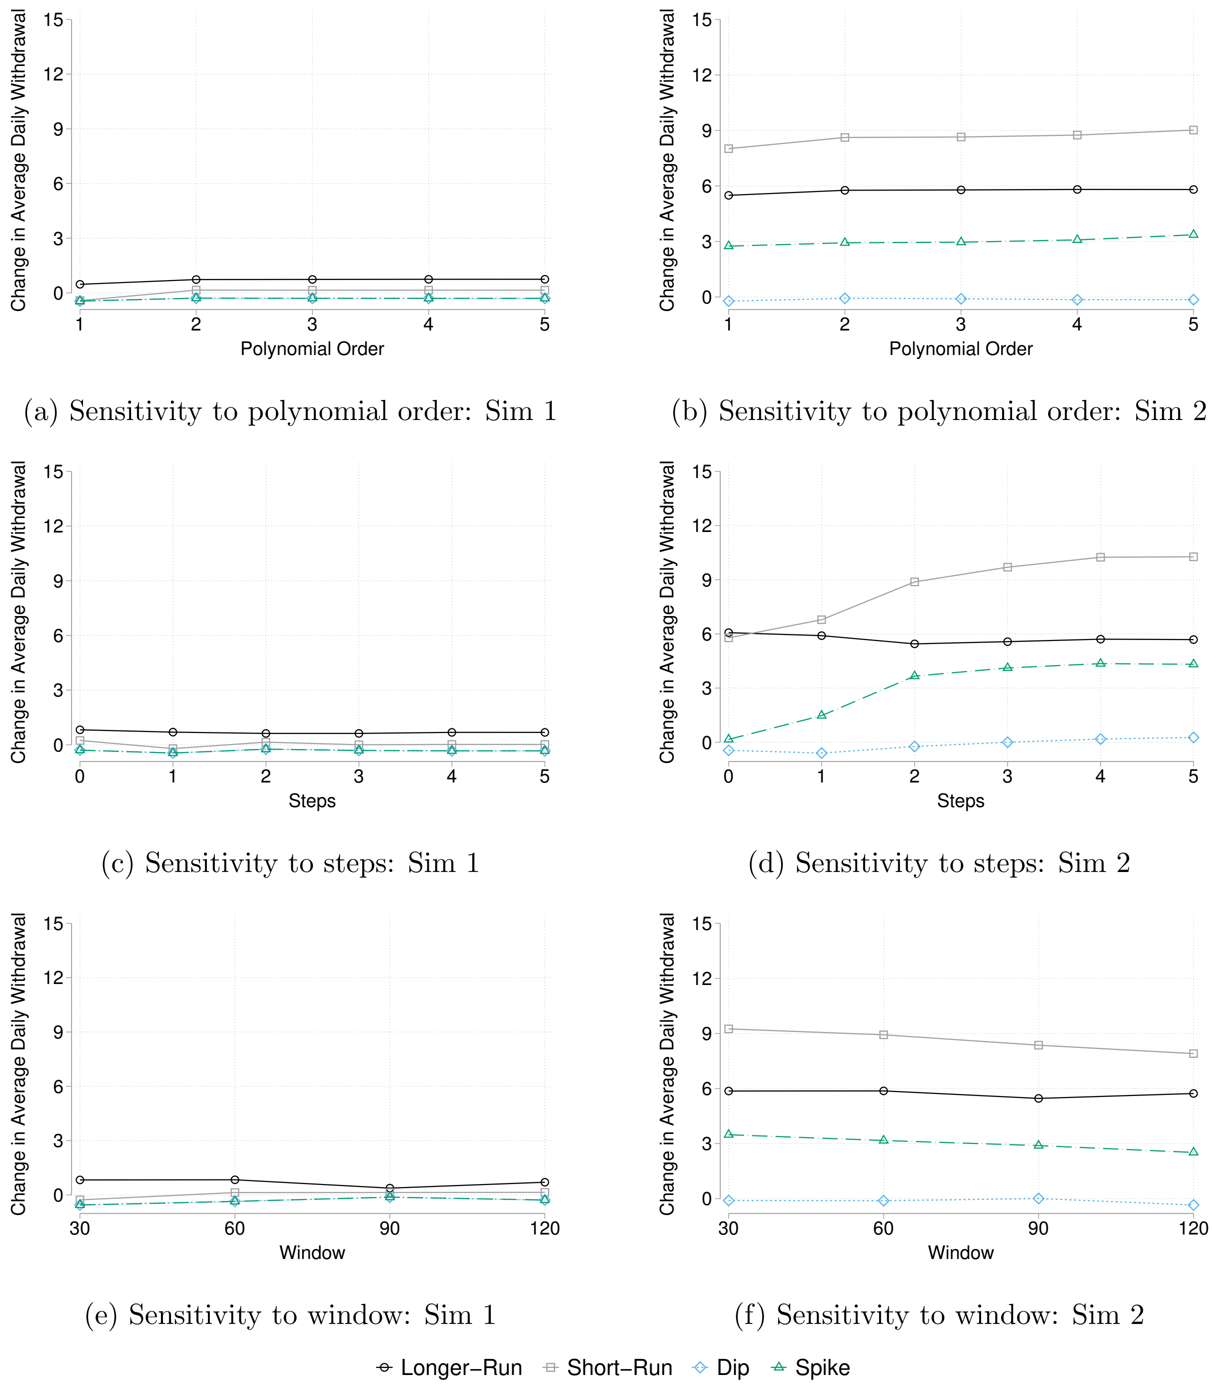

(a) Sensitivity to polynomial order: Sim 1

(b) Sensitivity to polynomial order: Sim 2

(c) Sensitivity to steps: Sim 1

(d) Sensitivity to steps: Sim 2

(e) Sensitivity to window: Sim 1

(f) Sensitivity to window: Sim 2

Notes: Figures show average value of Short-Run effect, Long-Run effect, Dip and Spike across different polynomial order, window and step assumptions. Each figure includes estimates from 120 different estimations.

Figure B.3: Simulated Data: Placebo and Permutation Tests

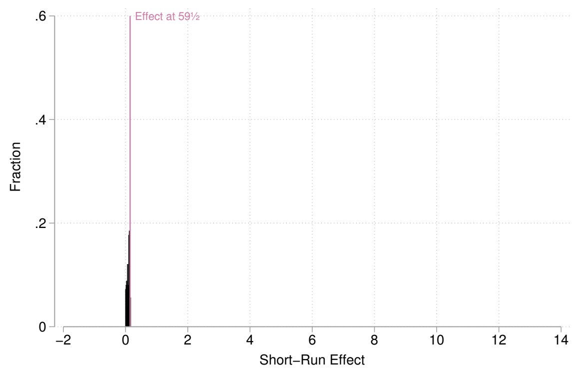

(a) Placebo: Sim 1

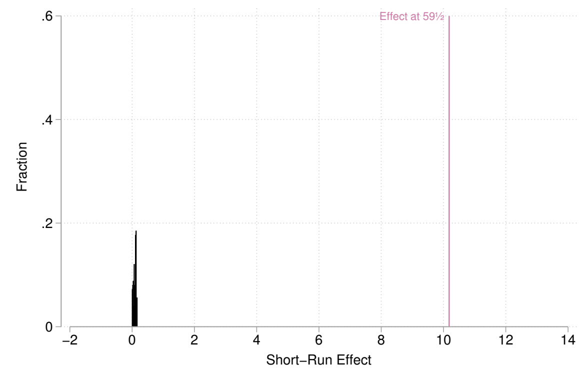

(b) Placebo: Sim 2

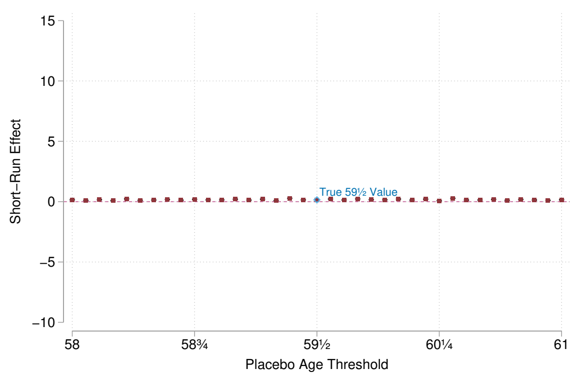

(c) Permutation: Sim 1

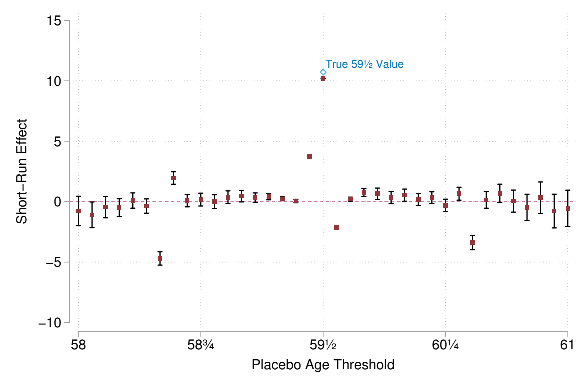

(d) Permutation: Sim 2

Notes: Top panels shows histogram of placebo values of Short-Run Effect across placebo age thresholds [57 and 9 months, 57 and 11 months] using data to the left of the age $59\frac{1}{2}$ threshold and [60 and 11 months, 61 and 1 month] using data to the right of the age $59\frac{1}{2}$ threshold for Sim 1 and Sim 2. Bottom panel shows estimated Short-Run Effect using full data and alternative age thresholds spanning [58 and 0 months, 61 and 0 months] for Sim 1 and Sim 2.

Appendix C: Additional Results

Table C.1: Changes in IRA Daily Withdrawals at Age $59\frac{1}{2}$ by Adjusted Gross Income and Fair Market Value Quartiles

	(1) Baseline Withdrawals	(2) Share Penalized	(3) Short-Run	(4) Long-Run	(5) Dip	(6) Spike
Adjusted Gross Income						
Bottom Quartile	3.15	0.139	12.86	5.44	-2.97	10.39
			(2.77)	(0.26)	(1.48)	(1.47)
Second Quartile	4.33	0.189	5.26	4.83	-3.69	4.12
			(2.41)	(0.22)	(1.29)	(1.28)
Third Quartile	5.33	0.291	6.11	4.95	-1.47	2.63
			(2.95)	(0.27)	(1.57)	(1.57)
Top Quartile	7.64	0.289	20.36	5.39	6.09	8.88
			(6.30)	(0.59)	(3.36)	(3.35)
Fair Market Value						
Bottom Quartile	1.63	0.469	2.96	1.03	1.69	0.24
			(1.34)	(0.12)	(0.71)	(0.71)
Second Quartile	1.56	0.42.4	3.91	3.25	0.09	0.56
			(1.65)	(0.15)	(0.88)	(0.88)
Third Quartile	2.22	0.287	8.87	5.40	0.25	3.22
			(2.35)	(0.22)	(1.26)	(1.25)
Top Quartile	15.60	0.152	30.08	12.41	-2.79	20.45
			(8.02)	(0.75)	(4.28)	(4.26)

SEs are in parentheses.

Notes: Standard errors calculated using the delta method. Short-Run and Long-Run Effects represent increase in average IRA daily withdrawals over 1 month before and after age $59\frac{1}{2}$, and from 2-3 months before age $59\frac{1}{2}$ to 2-3 months after age $59\frac{1}{2}$, respectively. Dip represents reduction in average IRA daily withdrawals from 2-3 months before age $59\frac{1}{2}$ to 0-1 month before age $59\frac{1}{2}$, and Spike represents increase in average IRA daily withdrawals from 2-3 months after age $59\frac{1}{2}$ to 0-1 month after age $59\frac{1}{2}$. Baseline represents average level of daily IRA withdrawals 2-3 months prior to age $59\frac{1}{2}$.

Appendix D: Sensitivity to parameters

Figure D.1: Sensitivity to Parameters

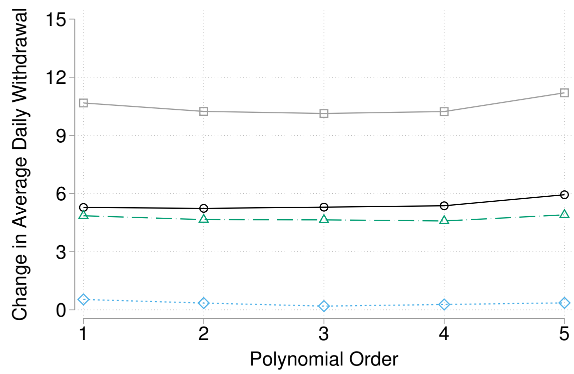

(a) Sensitivity to polynomial order

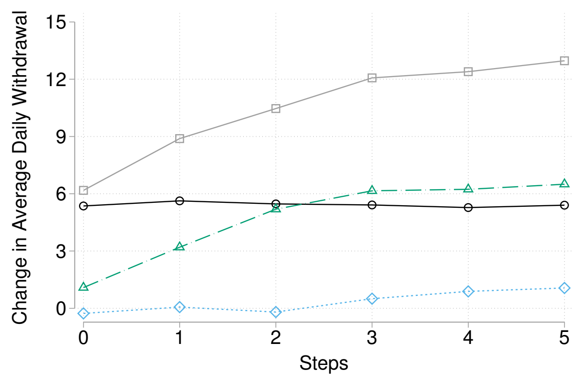

(b) Sensitivity to steps

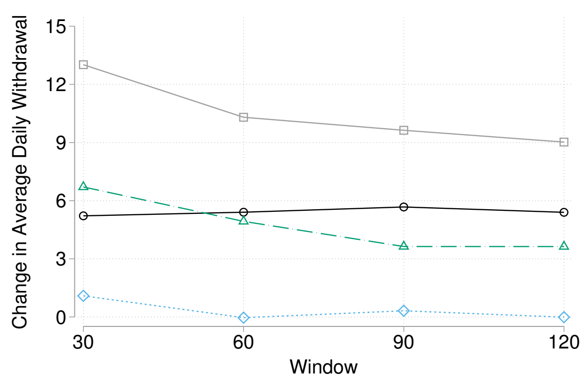

(c) Sensitivity to window

Longer−Run Short−Run Dip Spike

Notes: Figures show average value of Short-Run effect, Long-Run effect, Dip and Spike across different polynomial order, window and step assumptions. Each figure includes estimates from 120 different estimations.

Appendix E: Derivations

E.1 Derivation of Annual Withdrawals

Summing equation (1) over the event days in a year, we derive an expression for total annual withdrawals:

$$
\begin{aligned}
y_{be} &\equiv \sum_{j\in\mathcal{J}(b,e)} \tilde{y}_{bj} \\
&= \sum_{d\in t(b,e)} \tilde{\lambda}_d + \sum_{j\in\mathcal{J}(b,e)} \Big[f_l(j) \cdot \mathbf{1}\{j<0\} - D\cdot \mathbf{1}\{-w \le j < 0\} \\
&\qquad + \left[L + f_r(j)\right] \cdot \mathbf{1}\{j \ge 0\} + S \cdot \mathbf{1}\{0 \le j \le w\} + \tilde{\varepsilon}_{bj} \Big] \\
&= \lambda_{t(b,e)} + \sum_{j\in\mathcal{J}(b,e)} \Big[\sum_p \beta_{l,p} j^p \cdot \mathbf{1}\{j<0\} - D\cdot \mathbf{1}\{-w \le j < 0\} \\
&\qquad + \left[L + \sum_p \beta_{r,p} j^p\right] \cdot \mathbf{1}\{j \ge 0\} + S \cdot \mathbf{1}\{0 \le j \le w\} \Big] + \varepsilon_{be} \\
&= \lambda_{t(b,e)} + \sum_p \beta_{l,p} \left[\sum_{j\in\mathcal{J}(b,e)} j^p \cdot \mathbf{1}\{j<0\}\right] - D \left[\sum_{j\in\mathcal{J}(b,e)} \mathbf{1}\{-w \le j < 0\}\right] \\
&\quad + L\left[\sum_{j\in\mathcal{J}(b,e)} \mathbf{1}\{j\ge 0\}\right] + \sum_p \beta_{r,p}\left[\sum_{j\in\mathcal{J}(b,e)} j^p \cdot \mathbf{1}\{j \ge 0\}\right] + S\left[\sum_{j\in\mathcal{J}(b,e)} \mathbf{1}\{0 \le j < w\}\right] + \varepsilon_{be}
\end{aligned}
\tag{E.1}
$$

Here $\lambda_{t(b,e)} \equiv \sum_{d\in t(b,e)} \tilde{\lambda}_d$ is a calendar-year fixed effect and $\varepsilon_{be} \equiv \sum_{j\in\mathcal{J}(b,e)} \tilde{\varepsilon}_{bj}$ is an annual, mean-zero error term. In the third line, we express the smooth functions as polynomials, i.e. $f(j) = \sum_p \beta_p j^p$.

E.2 Value of Tax Deferral

We will quantify the tax implications of withdrawing early at time 0, relative to withdrawing at a later, counterfactual date t^*, after the penalty is permanently removed. The benefit to waiting is referred to as the value of "tax deferral." We assume the agent has access to a savings account with a return of r_t in each period between 0 and t^*. The amount to potentially be withdrawn is $\triangle With$. Suppose there are no taxes. We define r as the implied total return over this period:

$$\begin{aligned} (1+r) &\equiv \prod_{t=1}^{t^*}(1+r_t) \\ \rightarrow r &\equiv \prod_{t=1}^{t^*}(1+r_t)-1 \end{aligned} \tag{E.2}$$

On the other hand, if the agent must pay income taxes τ_t on the return in each period, then the after-tax return is smaller. We define τ_r as the implied decrease in total returns:

$$\begin{aligned} (1+r(1-\tau_r)) \equiv & \quad \prod_{t=1}^{t^*}(1+r_t(1-\tau_t)) \\ \rightarrow \tau_r \equiv & \quad 1-\frac{\prod_{t=1}^{t^*}(1+r_t(1-\tau_t))-1}{r} \\ = & \quad \frac{\prod_{t=1}^{t^*}(1+r_t)-\prod_{t=1}^{t^*}(1+r_t(1-\tau_t))}{\prod_{t=1}^{t^*}(1+r_t)-1} \end{aligned} \tag{E.3}$$

Using the above identities, note the difference in returns on savings between the two cases:

$$\prod_{t=1}^{t^*}(1+r_t)-\prod_{t=1}^{t^*}(1+r_t(1-\tau_t))=(1+r)-(1+r(1-\tau_r))=\tau_r r \tag{E.4}$$

This quantifies the tax savings afforded to the agent who saves in a tax deferred account instead of a normal savings account.

To quantify the value of tax deferral, we calculate effective taxes on savings under two scenarios: an early withdrawal in period 0 and a delayed withdrawal in period t^*. In either case, we contrast the amount of wealth available in period t^* in the absence of any taxes and the amount of wealth available in period t^* after taxes. The difference between the two is the effective tax on savings, as valued in period t^*.

If an agent makes an early withdrawal, a tax of τ_0 must be paid on the withdrawal amount. That after-tax principal could then be saved and earn a return of r_t in each subsequent period. The return to these savings will be taxed in each period at τ_t. If the agent holds the amount in a tax-deferred account, i.e. an IRA, no taxes are paid today on the principal, and taxes are not collected on interest income in each period. Upon withdrawal in period t^*, a tax of τ_{t^*} is paid on the principal and interest earned. We can therefore quantify the value of tax deferral as follows:

$$\begin{aligned}
\text{Tax Deferral} &= \underbrace{\Bigg[\underbrace{\triangle With \cdot (1+r)}_{\text{No tax withdrawal}} - \underbrace{\triangle With \cdot (1-\tau_0)(1+r(1-\tau_r))}_{\text{After tax withdrawal}}\Bigg]}_{\text{Effective tax w/ early withdrawal}} \\
&\quad - \underbrace{\Bigg[\underbrace{\triangle With \cdot (1+r)}_{\text{No tax withdrawal}} - \underbrace{\triangle With \cdot (1+r)(1-\tau_{t^*})}_{\text{After tax withdrawal}}\Bigg]}_{\text{Effective tax w/ delayed withdrawal}} \\
&= \Big[(1-\tau_{t^*})(1+r) - (1-\tau_0)(1+r(1-\tau_r))\Big] \times \triangle With \\
&\approx \underbrace{\tau_r r}_{\text{Avoided taxes on Interest}} \times \underbrace{(1-\tau)\triangle With}_{\text{After-tax principal}}
\end{aligned} \tag{E.5}$$

In the last line, we've made a simplification that the tax rate is constant throughout, i.e. $\tau_0 = \tau_{t^*} = \tau$. Though unrealistic, this allows us to focus on the primary benefit of an IRA, which is the deferral of when taxes are paid. The key term $\tau_r r$ is the forgone tax on interest that we derived above, which is applied to the after-tax principal, $(1-\tau)\triangle With$. Intuitively, the principal investment is taxed under either scheme, but this after-tax principal earns a higher return in a tax-deferred account. This expression represents the value of the tax deferral in period t^*. To get the present value of that tax deferral, we would have to discount future payments to present value using an appropriate discount factor δ.