Research Dialogue | Issue no. 136
June 2017

TIAA Institute

How Do Distributions From Retirement Accounts Respond to Early Withdrawal Penalties?

Evidence From Administrative Tax Returns

Gopi Shah Goda,
Stanford University, NBER,
TIAA Institute Fellow

Damon Jones,
University of Chicago,
NBER

Shanthi Ramnath,
U.S. Treasury

Abstract

The design of retirement savings accounts must balance the long-term goal of retirement wealth accrual with the potential need for liquidity in the short-term. Penalties on pre-retirement withdrawals provide a possible lever for striking this balance. In the United States, penalties are typically limited to 10 percent of withdrawn funds and several exceptions are available in order to provide access to savings in response to a shock. In this paper, we investigate how individuals respond to the removal of the 10 percent penalty imposed on Individual Retirement Account (IRA) withdrawals prior to the account holder turning 59 ½. Our analysis employs rich tax records from the Internal Revenue Service (IRS) and develops new empirical techniques which allow us to use annual data to better understand movements at higher levels of frequency. Our findings show evidence of a 93 percent increase in annual withdrawals on average among our population, and suggests that much of this increase in withdrawals is coming from new people making withdrawals once they cross the age 59 ½ threshold.

Email: gopi@stanford.edu; damonjones@uchicago.edu; shanthi.ramnath@treasury.gov. We would like to thank 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 findings and conclusions expressed are solely those of the authors and do not represent the views of SSA, any agency of the Federal Government, the TIAA Institute, the Pension Research Council, or the NBER.

1. Introduction

In the United States, Americans have an estimated $14.4 trillion invested in employer- sponsored defined contribution plans and individual retirement accounts (Investment Company Institute, 2015). These funds typically receive preferential tax treatment, which allows households to accumulate retirement savings at a faster rate than in normal savings vehicles. Furthermore, the accounts are relatively illiquid, as a penalty is typically imposed for withdrawals occurring before the account holder turns 59 ½. This penalty is designed to dissuade people from accessing these funds prior to retirement and may result in more retirement savings than would otherwise be achieved. On the other hand, this illiquidity may limit the ability of households to smooth consumption prior to retirement.

There are several avenues to partially or completely liquidate funds in tax-preferred retirement savings accounts prior to retirement. First, many accounts grant exceptions from the penalty for several reasons including death or disability, education expenses, first-time home purchases, and unreimbursed medical expenses. In addition, job transitions can provide opportunities to liquidate tax-preferred retirement savings accounts with funds less than a specified threshold, and some accounts allow loans which may become distributions if not paid back upon job separation.[1] Finally, many accounts allow distributions to be taken for any reason subject to a penalty being paid.

Recent evidence suggests that these pre-retirement withdrawals, known as "leakage," are substantial, amounting to approximately $0.40 of every $1 contributed into the account prior to the age of 55 (Bryant, Holden and Sabelhaus, 2010; Argento, Bryant and Sabelhaus, 2015). Leakage reduces wealth available for retirement substantially, and the potential to access retirement funds prior to retirement could lead present-biased individuals to accumulate lower levels of retirement wealth (Beshears et al., 2014, 2015*a*; Goda et al., 2015). This may suggest that there is a benefit to increasing the penalty and making the account a more effective commitment device.

On the other hand, wealth accumulated in retirement savings accounts can also provide an important form of insurance. Indeed, previous studies find that early withdrawals are strongly correlated with shocks to income or marital status or represent consumption- smoothing behavior by liquidity-constrained households who experience financial shocks (Amromin and Smith, 2003; Argento, Bryant and Sabelhaus, 2015). If retirement savings accounts allow the ability to insure against negative consumption shocks, then some level of liquidity prior to retirement may be optimal. In addition, offering this liquidity may make it more attractive to save in retirement savings accounts relative to accounts where pre-retirement withdrawals are forbidden.

The potential consumption-smoothing benefits retirement savings accounts can provide may be at odds with the goals of retirement wealth accumulation. As a result, there has been recent discussion regarding adjusting the age threshold for penalty-free withdrawals (Munnell and Webb, 2015) or changing the amount of the penalty (Beshears et al., 2014). Moreover, several other developed countries, which generally lack options for early withdrawal, are in the process of discussing providing early access to retirement savings (Beshears et al., 2015*a*; Agarwal, Pan and Qian, 2016). Despite these active policy debates, there is not a large amount of literature seeking to understand the implications of these potential policies.

In this paper, we examine the withdrawal behavior of individuals as they cross the age 59 ½ threshold in retirement savings accounts when the penalty for early withdrawals is removed. 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 which contain information regarding individuals' retirement accounts, contributions, distributions, as well as one's filing status, adjusted gross income, wages, and other items collected by tax forms. While these data have several advantages, the fact that they can only obtained on an annual basis rather than higher levels of frequency means that it is difficult to disentangle general increases in retirement distributions as individuals age from increases occurring as a result of the removal of the penalty at age 59 ½.

1. For instance, the IRS waives any penalties for workers aged 55 and older after a job termination.

In order to identify the response in retirement account withdrawals, we exploit differences in exposure to penalty-free withdrawal within a calendar year stemming from variation in one's date of birth. For instance, someone whose birthday is July 1, 1949 attains 59 ½ on January 1, 2009 and thus has a full year of exposure to penalty-free withdrawals in 2009. By contrast, an individual born on June 30, 1950 attains age 59 ½ at the end of the year on December 30, 2009 and only experiences one day of penalty-free withdrawal in 2009. Building on that intuition, we introduce a novel method for using annual data to parametrically recover an event study at age 59 ½.

We find that increases in annual distributions in the calendar year one turns 59 ½, relative to the previous calendar year, are larger for individuals who attain age 59 ½ early in the year relative to those who attain age 59 ½ late in the year. Our empirical estimates indicate that crossing the age 59 ½ threshold leads to an approximately $1,600 increase in annual distributions from Individual Retirement Accounts (IRAs), on average. Our results suggest that this increase is primarily due to additional people taking withdrawals after the penalty is lifted rather than higher distributions among those who were withdrawing prior to age 59 ½ .

Our paper builds on related literature that examines how withdrawals from retirement savings accounts change over the lifecycle and in response to various provisions. Perhaps most relevant, recent work by Agarwal, Pan and Qian (2016) examines how withdrawals from pension savings in Singapore responds to a sharp change in the ability to cash out savings at age 55. Using data from a large bank, the authors construct a monthly event study surrounding age 55 and show that account balances and credit card spending increase upon turning 55, while credit card debt decreases. Prior work using U.S. data show increases in withdrawals by age (e.g., Sabelhaus (2000)), but does not allow for higher-frequency event studies to uncover the relationship between withdrawal penalties and distribution rates.

Recent studies examine withdrawals behavior surrounding the age threshold for required minimum distributions. Poterba, Venti and Wise (2013) find that withdrawal behavior increases sharply after age 70 ½ using data from the SIPP, suggesting that households tend to preserve retirement assets to self-insure against large and uncertain late-life expenses. Brown, Poterba and Richardson (2014) examine how the 2009 one-time suspension of the rules associated with required minimum distributions affected distributions for TIAA participants and find that one third of those affected by the rules discontinued their distributions when the rules were suspended. Using administrative tax data, Mortenson, Schramm and Whitten (2016) similarly find that required minimum distributions cause funds to be drawn down more quickly than otherwise, and, additionally, that some accounts are closed in response to the policy.

We make several contributions to this literature. First, we provide, to our knowledge, the first causal estimates of the effect of removing the 10 percent penalty from pre-retirement withdrawals on withdrawal behavior in the U.S. Under the assumption that other characteristics that affect distribution behavior vary smoothly across the age 59 ½ threshold, our estimates can be interpreted as the result of the change in penalty rather than other factors.

Second, we use a novel, comprehensive data source that provides high-quality data on distributions from information returns provided by the IRS. Given the relatively small numbers of individuals taking withdrawals from retirement savings accounts near the age 59 ½ threshold, household surveys are unlikely to uncover any changes occurring precisely at age 59 ½. In addition, household surveys may underreport distributions from retirement savings accounts, as even distributions recorded on Form 1040 are approximately 20 percent lower than implied by information returns (Argento, Bryant and Sabelhaus, 2015).

Finally, we develop empirical techniques to convert data at a lower frequency into a higher frequency event study by exploiting variation by date of birth. These techniques are similar to, but distinct from, techniques that exploit differences in the distribution of temperature each year to identify the effect of particular daily temperatures on outcomes in the climate change literature (Deschênes and Greenstone, 2011; Deryugina and Hsiang, 2014). Our method can potentially be used in a variety of different settings, including, for example, understanding the effect of sharp changes in eligibility for Social Security on related outcomes.

The remainder of the paper proceeds as follows. Section 2 describes institutional features and the data we use for the study, and Section 3 lays out our empirical strategy. We discuss results in Section 4 and conclude in Section 5.

2. Background and data

2.1 Retirement liquidity in the U.S.

A large component of retirement savings in the U.S. is in tax-preferred savings accounts, including both employer-sponsored defined contribution plans (e.g., 401(k)s) and Individual Retirement Accounts (IRAs). These accounts allow individuals to contribute funds annually, up to a set maximum. Contributions are either made with pre-tax assets and taxed when withdrawn, as in the case of Traditional IRAs or 401(k)s, or made with after-tax assets and exempt from taxes when withdrawn, as in the case of Roth IRAs or Roth 401(k)s.

In order to encourage individuals to use the proceeds from these accounts for retirement, the government imposes various restrictions or penalties against withdrawing funds for other purposes. The restriction depends on precisely which type of account is being withdrawn from. Typically, traditional IRAs allow early withdrawals for any reason, but these early withdrawals are subject to a 10 percent penalty. Exceptions to the penalty are made in the event of death or disability, for first-time homebuyers, education expenses, health insurance premiums while unemployed, and unreimbursed medical expenses. Since Roth IRA contributions are made on an after-tax basis, withdrawing the basis—and not the earnings—can be done without penalty.

Pre-retirement distributions from 401(k) plans can be made only in the event of a hard- ship, or an immediate and heavy financial need. Certain expenses are deemed to be immediate and heavy, including certain 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).

All penalties and restrictions are lifted once an individual turns 59 ½. The IRS calculates age 59 ½ 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. 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 ½. Since this day does not exist, the rules indicate that one would calculate the residual days left over at the end of the month (three in this case, since February ends on February 28) and advance that many days forward (March 3, 2030 in this example). Individuals born on leap days turn 59 ½ on September 1 in the year in which they turn 59. Put another way, the IRS considers them age 59 ½ on the same day as someone born on March 1 in the same year.

While not the focus of this paper, there is also a sharp change in rules regarding withdrawals when an individual turns 70 ½ and is subject to required minimum distributions (RMDs). RMDs apply to all employer-sponsored retirement plans and traditional IRAs and specify minimum amounts that an account owner must withdraw annually starting the year he or she attains age 70 ½.[2] These rules are designed to limit the amount of tax deferral provided to retirement savings accounts.[3]

It is worth noting that the liquidity in retirement savings accounts in the U.S. is generally higher than other developed countries. Beshears et al. (2015*b*) compare the liquidity in retirement savings systems across six developed countries and show that the U.S. has a much more liquid system with relatively low penalties for early withdrawals, and several exceptions for penalty-free withdrawals.

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 for tax years 1999 through 2013 who are alive in the year they turn 57 ½. This sample restriction ensures that our data contain tax years two years before and after each individual turns 59 ½. Our dataset contains information on household income (Form 1040), wage earnings and employee contributions to employer- sponsored retirement plans (Form W2), distributions from IRAs and employer-sponsored retirement plans (Form 1099R), contributions to and account balances of IRAs (Form 5498), and tax amounts on early distributions (Form 5329). Because the data are unedited,

2. For employer-sponsored retirement plans, individuals are exempt from RMDs if they are not retired.
3. Note that Roth IRAs do not require minimum distributions until after the death of the account holder.

we make a number of restrictions in an effort to remove observations with erroneous information. We drop roughly 1.5 million observations due to death and birthdates that do not exist (e.g. September 31).

Our analysis focuses on distributions from IRAs due to some important data limitations. First, unlike Form 5498 which provides the fair market value of an IRA annually, there is no tax form at the individual level that reports account balances for defined contribution plans. This makes it difficult to select a sample of individuals who are at risk of withdrawing funds from these accounts.[4] Second, while distributions from defined contribution plans are reported on 1099-R forms, they are undistinguishable from defined benefit payments. By contrast, IRA distributions can be separately identified due to a checkbox on the 1099-R tax form.

As described in the previous section, the penalties differ somewhat for 401(k)s and IRAs, as 401(k) plans only allow hardship withdrawals prior to age 59 ½ while IRAs allow withdrawals for any reason. 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 ½, since many individuals roll over their employer-sponsored retirement accounts into IRAs prior to retirement.

Our main analysis sample contains individuals who have a positive fair market value in at least one IRA as reported on Form 5498 in the year they turn 57 ½. While our data are at the individual level, we collapse the data by individual date of birth to perform our analysis, which exploits variation in exposure to the penalty-free withdrawal period using variation in date of birth. Therefore, our total number of observations is 14,608 date-of-birth-by-year cells, representing 12,445,087 individuals or 36% percent of the population who attains age 57 ½ in our analysis period.

Table 1 contains descriptive statistics on our sample. The data represent information from tax years in which the age in the column heading is attained. Just under half of our sample is male and almost three quarters file a joint return. The average adjusted gross income in our sample is $113,821. This value is relatively high both because of our sample restriction that eliminates those without assets in IRA accounts and the older ages in our sample. The fraction of our sample that takes distributions from their Traditional IRA is 7 or 8 percent prior to the tax year in which the individual turns age 59 ½, then increases to 13 percent and 16 percent in the following two years. The amount withdrawn conditional on taking a distribution is approximately $21,000 annually. The fact that this amount does not vary markedly around the age 59 ½ threshold suggests that any increase in average withdrawals occurring at the age 59 ½ threshold may be more likely to be on the extensive margin. Importantly, these simple comparisons of annual distributions across ages do not allow identification of responses to the sharp reduction in the penalty occurring when one attains age 59 ½ as these increases could simply represent increasing shares of individuals taking distributions as they get closer to retirement.

3. Empirical strategy

While we observe distributions from IRA accounts at an annual frequency, we would like to understand how distributions evolve at a more granular level and, in particular, in a neighborhood near age 59 ½. Our empirical strategy leverages variation in exposure to early withdrawal penalties driven by date of birth to recover patterns at a subannual frequency. For example, take two individuals, one born on June 30, 1950 and another born on July 1, 1950. According to IRS rules, the former turns 59 ½ on December 30th, 2009, while the latter turns 59 ½ on January 1st, 2010. Differences in their annual distributions in 2009 can be related to the fact that one person has experienced two days of penalty free distributions while the other faced the penalty the entire year. We generalize this notion below. First, we present a method, relying on strong parametric assumptions, that uses annual patterns in year-to-year distribution levels to test for a discontinuous effect of the age 59 ½ threshold. Second, we provide a less restrictive approach that allows us to estimate an event study in retirement distributions at age 59 ½, at subannual frequencies: i.e. quarterly, monthly, weekly, and daily.

4. The tax data do contain an indicator of whether one's current employer offers a defined contribution plan, and data on contributions made to defined contribution plans; however, both of these are noisy indicators of individuals with a positive balance.

We motivate our empirical approach with a model of average daily retirement distributions. We assume that distributions would evolve in a continuous and gradual fashion from day-to-day, in the absence of sharp changes in withdrawal penalties. Suppose the daily pattern of distributions can be characterized as follows:

$$y_{bd} = \tilde{\alpha} + \tilde{\lambda}_d + f\left(d - b - a^*\right) + D \cdot \mathbf{1}\left\{d - b \geq a^*\right\} + \varepsilon_{bd}, \quad (1)$$

where y_{bd} is the average daily distribution on day d among individuals born on day b, $\tilde{\lambda}_d$ is a calendar day fixed effect, a^* is the number of days it takes to reach age 59 ½, and ε_{bd} is a mean-zero error term. The function $f(\cdot)$ governs the age pattern of retirement distributions, and its argument is measured relative to age 59 ½. The function $1\{\cdot\}$ is an indicator function, and, thus, the parameter D represents an additively separable shift in average retirement distributions upon reaching age 59 ½.

3.1 Annual patterns

As a first step toward testing for a discontinuous change in behavior upon turning age 59 ½, we show what can be inferred from annual retirement distribution patterns. We make the extreme assumption that, aside from the possible discontinuity at age 59 ½, retirement distributions are locally linear in age, i.e. $f(j) \equiv c \cdot j$. Let e measure event time in years. That is, $e = 0$ in the year in which one reaches age 59 ½, $e = -1$ in the year in which one turns 58 ½, and so forth. We define the average annual distribution during event year e, for an individual born on day b, as $y_{be} \equiv \sum y_{bd}$.

Now, suppose we group individuals into cells based on the quarter in which age 59 ½ is reached and event year. Within this cell, we will calculate average annual retirement distributions, denoted $\overline{y}_{qe} \equiv \sum_{b:q(b)=q} (N_b \cdot y_{be}) \Big/ \sum_{b:q(b)=q} N_b$, where N_b is the number of individuals born on day b and the mapping $q(b) \in \{1, 2, 3, 4\}$ returns the quarter in which someone born on day b reaches age 59 ½. Finally, let the change in this average from event year $e-1$ to e be denoted as $\triangle y_{qe} \equiv \overline{y}_{qe} - \overline{y}_{qe-1}$. Using equation 1, we can show the following:

$$\begin{aligned} \triangle y_{q,-1} - \triangle y_{q',-1} &\approx 0 \\ \triangle y_{q,0} - \triangle y_{q',0} &\approx [q' - q] \cdot (365/4) \cdot D \\ \triangle y_{q,1} - \triangle y_{q',1} &\approx -[q' - q] \cdot (365/4) \cdot D. \end{aligned} \quad (2)$$

In words, we first measure the change in retirement distributions from the year in which one turns 57 ½ to the year in which one turns 58 ½. The difference in this change across different quarters of reaching age 59 ½ is approximately zero. Second, we measure the change in distributions from the year in which one turns 58 ½ to the year in which one turns 59 ½. The difference in this change across different quarters of reaching age 59 ½ is approximately proportional to the difference in quarters. Finally, when comparing the change in distributions from the year in which one turns 59 ½ to the year in which one turns 60 ½, the difference in average annual distributions decreases approximately linearly in the difference in quarters of reaching 59 ½. We can test the predictions by estimating the following regression for different values of e:

$$\triangle y_{qe} = \alpha + D_e \cdot (-365/4)\, q + \varepsilon_{qe}, \quad (3)$$

where $e \in \{-1, 0, 1\}$. Under the linearity assumption, we expect $\hat{D}_{-1} = 0$ and $\hat{D}_0 = -\hat{D}_1$.

Likewise, if we group individuals into cells by month of reaching age 59 ½ and event year, and calculate the change in average annual retirement distributions, similarly denoted $\triangle y_{me}$, we have:

$$\begin{aligned} \triangle y_{m,-1} - \triangle y_{m',-1} &\approx 0 \\ \triangle y_{m,0} - \triangle y_{m',0} &\approx [m' - m] \cdot (365/12) \cdot D \\ \triangle y_{m,1} - \triangle y_{m',1} &\approx -[m' - m] \cdot (365/12) \cdot D. \end{aligned} \quad (4)$$

These admit a similar regression using data grouped by month of birth:

$$\triangle y_{me} = \alpha + D_e \cdot (-365/12)\, m + \varepsilon_{me}, \quad (5)$$

with similar predictions for $\hat{D}_e$ as in the case of quarterly averages.

Figure 1 illustrates the intuition behind these results. Panel A demonstrates the pattern of distributions that would arise in the case where $D = 0$ in equation 1. The horizontal axis measures age, and the calendar years in which one turns 58 ½ and 59 ½ are highlighted. The vertical axis measures the average daily distribution. The drawing on the left depicts an "older" agent, who turns 59 ½ relatively early in the year.

Note, at age 59 ½ there is no discrete jump in distributions. The area under the curve aggregates daily distributions into annual amounts, and the area shaded in light red measures the change in average annual distributions from the year in which one reaches 58 ½ to the year in which one reaches 59 ½, i.e. $\triangle y_{b,0}$. Similarly, the drawing on the left depicts the same patterns for a younger agent, who reaches 59 ½ later in the calendar year. Although the level of the annual distributions differs between the two, the *change* in the annual distribution is the same.

Alternatively, Panel B of Figure 1 presents the pattern of distributions in the case where $D > 0$. Here, we see that at age 59 ½, there is an upward shift in the average daily distributions. Furthermore, when comparing the change in distributions, we now see that the "older" agent, who experiences a longer time without early withdrawal penalties, exhibits a larger change distributions.

Although these results rely on strong functional form assumptions, they deliver sharp predictions regarding year-to-year changes in annual distributions across different quarters and months of reaching 59 ½. In particular, the above results imply that increases in annual distributions are roughly constant across quarter and/or month of reaching 59 ½ between event years -2 and -1, are monotonically increasing in quarter and month of reaching 59 ½ between event years -1 and 0, and are monotonically decreasing between event years 0 and 1. In the next section, we develop a more flexible approach to learning about the shape of the $f(\cdot)$ function and the parameter D.

3.2 Estimated daily event study

Building upon the intuition in the previous section, we now relax the assumptions made about the functional form of $f(\cdot)$ and instead estimate this function using a flexible polynomial. To better parallel the structure of our tax data, we will shift time relative to age a^* and collapse data to an annual level. Let j measure age in event time, i.e. age relative to the date on which one turns 59 ½. Formally, let $j \equiv d - b - a^*$. Let the mapping $t(b,e)$ be the year in which someone born on day b reaches year e in event time. For ex- ample, $t(b = 10June1950, e = 0) = 2009$. Likewise, the mappings $\underline{d}(b,e)$ and $\overline{d}(b,e)$ are the calendar dates for January 1 and December 31 in the year $t(b,e)$.

We can now express annual retirement distributions as follows:

$$\begin{aligned} y_{be} &\equiv \sum_{d=\underline{d}(b,e)}^{\overline{d}(b,e)} y_{bd} \\ &= \sum_{d=\underline{d}(b,e)}^{\overline{d}(b,e)} \tilde{\alpha} + \sum_{d=\underline{d}(b,e)}^{\overline{d}(b,e)} \tilde{\lambda}_d + \sum_{j=\underline{d}(b,e)-b-a^*}^{\overline{d}(b,e)-b-a^*} f(j) + \sum_{j=\underline{d}(b,e)-b-a^*}^{\overline{d}(b,e)-b-a^*} D \cdot \mathbf{1}\{j \geq 0\} + \sum_{d=\underline{d}(b,e)}^{\overline{d}(b,e)} \varepsilon_{bd} \\ &= \alpha + \lambda_{t(b,e)} + \sum_{j=\underline{d}(b,e)-b-a^*}^{\overline{d}(b,e)-b-a^*} [f(j) + D \cdot \mathbf{1}\{j \geq 0\}] + \varepsilon_{be}, \end{aligned} \tag{6}$$

where $\alpha \equiv 365 \cdot \tilde{\alpha}$ is a constant, $\lambda_t \equiv L_t \cdot \tilde{\alpha} + \sum_{d=1Jan,t}^{31Dec,t} \tilde{\lambda}_d$ is a calendar year fixed effect, L_t is an indicator for a leap year, and ε_{be} is a mean-zero error term. We fit $f(\cdot)$ with a flexible polynomial, using the specification in equation 6. In particular, we use polynomials of order one, three, and five, and additionally allow the coefficients to differ on either side of age 59 ½. Our key parameter of interest is D, which captures any sharp change in retirement distributions upon turning age 59 ½. The method can also be adapted to model average distributions at lower frequencies, i.e. weekly, monthly, or quarterly.

3.3 Simulations

In Appendix A we demonstrate our method using simulated data. We simulate 10 cohorts of individuals, each with four years of daily distributions, drawn to match key moments from the actual annual retirement distributions in our data. Figure A.1 shows the simulated pattern of daily distributions two years before, and one year following age 59 ½. We model a discrete jump in daily distributions of $10 once an individual no longer faces early withdrawal penalties. We also introduce a limited amount of curvature away from the threshold. We then collapse the data to annual frequencies, as is observed in our tax data.

We show in Figure A.2 patterns in annual retirement distributions by quarter and month of birth. As can be seen, the predictions in Section 3.1 are largely reflected in the simulated data. The increase in distributions from year to year is related the difference in exposure to penalty-free withdrawal opportunities. Next, in Figure A.3, we apply both our parametric estimator of the event study. As can be seen, we are able to closely recreate the true, underlying pattern for daily distributions. In Table A.1 we report the results from the regressions in equations 3 and 5. The estimates of D using either D_0 or D_1 at the quarterly or monthly frequency

are very close to the true value of $10. The parameter D_{-1} does not exactly equal zero, owing to the fact that we do not use a linear functional form for $f(\cdot)$ in our simulations. However, it's value is economically insignificant and an order of magnitude smaller than the other estimates. Table A.2 shows that when we use our more generalized approach, our point estimates of the jump in withdrawals at age 59 ½ closely match the true value used in the simulates, $10, albeit with some attenuation for the most coarse specification of quarterly aggregation. This is not surprising, as individuals are on average only exposed to penalty-free withdrawal for half of the quarter in which they turn 59 ½.

4. Results

4.1 Annual patterns

We first investigate how annual distributions vary in calendar years in which individuals attain ages 57 ½, 58 ½, 59 ½, and 60 ½ based on the exposure to time with penalty-free withdrawals. The exposure to time with penalty-free withdrawals depends on one's quarter or month of birth, as discussed in Section 3.1. Figure 2 shows annual distributions in different calendar years by the penalty-free exposure period. The top panel groups individuals by birth quarter, while the bottom panel organizes the sample by month of birth. Individuals represented by the line corresponding to 4 months of penalty-free withdrawal, for example, include those who turn 59 ½ between August 1 and September 1 (i.e., birthdays in February).

As shown in Figure 2, all of the groups have similar annual distributions in the years they turn 57 ½ and 58 ½. However, the lines begin to diverge in the year they turn 59 ½. In particular, those who have more months of penalty-free withdrawal in the calendar year in which they turn 59 ½ also have larger increases in their annual distributions in that calendar year. Note, this pattern is consistent with Panel B of Figure 1 as well. The spread in the average distributions by months of penalty-free withdrawal shrinks considerably in the year individuals attain 60 ½, suggesting that the higher rate of withdrawals persists as individuals age, but not differentially across birth months.

These figures provide evidence that the removal of the penalty at age 59 ½ is driving the patterns seen in the data and largely track the predictions made in Section 3.1, where we assume a linear relationship between distributions and age. In particular, the increase in distributions when moving from age 58 ½ to 59 ½ is robustly monotonic in quarters or months of exposure to penalty-free withdrawals, as is predicted in the approximations in 2 and 4. To see what magnitude of increase in distributions at age 59 ½ is implied by the figure, we estimate equations 3 and 5 in Table 2. Assuming a linear functional form for $f(\cdot)$, we estimate an increase of between $4 and $5 in distributions (i.e. from our estimates of D_0 and D_1). Our estimates of D_{-1}, however, reject a purely linear relationship between age and distributions. Nonetheless, the patterns here are largely consistent with a discontinuous increase in distributions upon reaching age 59 ½.

4.2 Estimated event study

We next perform an event study analysis in order to trace out daily withdrawal rates before and after individuals turn 59 ½. The results are from our fully-parametric analysis based on equation 1, where $f(\cdot)$ is modeled using a linear, cubic, or quintic polynomial and the analysis is carried out at the daily level are depicted in Figure 3. The vertical axes in the figures represent average daily withdrawal amounts, and we show both the point estimates and 95 percent confidence intervals. All of the fully-parametric specifications show evidence of a break at event time 0, which represents the day that individuals attain age 59 ½ and can begin making withdrawals from their IRA without any penalty. The higher-order polynomials reveal some of the underlying features of the data to a richer extent than the linear functional form. In all cases, daily withdrawals appear to be largely flat prior to age 59 ½. Thereafter, appears to be a spike in distributions at age 59 ½, followed by a decrease in withdrawals to a new level higher than prior ages. In Appendix B, we include similar figures using different levels of aggregation under the linear functional form. The vertical axes in these figures again represent daily withdrawal amounts, and show similar patterns.[5]

Table 3 reports the results from regressions that correspond to the event study analysis using different levels of aggregation and functional form assumptions. The regression equation, in the case of a daily frequency, is shown in equation 1 and includes calendar year fixed

5. Additional figures with higher-order polynomials under these alternative frequencies are omitted in the interest of space and are available from the authors upon request.

effects. The quarterly, monthly, and weekly frequencies are estimated from parallel specifications. The coefficient that is reported corresponds to the parameter D, and represents the sharp increase in daily withdrawals occurring when crossing the age 59 ½ threshold. We also report standard errors below the estimated coefficient.

The results show that using a daily level of aggregation and a linear functional form, lifting the 10 percent penalty on withdrawals leads to an increase in the daily withdrawal rate of $4.32, or $1,576.80 annually. This coefficient is precisely estimated and statistically different from zero. The other reported coefficients in the table similarly show strong evidence that the removal of the penalty significantly increases withdrawals from IRAs. The range of estimates varies between $3 and $8.50.

The increase in the rate of withdrawals could arise either from the extensive margin, if a larger share of individuals are taking distributions, the intensive margin, if the share of individuals accessing their IRAs does not change but the average amount conditional on taking distributions increases, or a combination of both. Relative to an unconditional average annual withdrawal rate of $1,682.36 ($21,029.47 × 0.08 from Table 1) in the calendar year individuals turn 59 ½, an increase of $1,576.80 represents roughly a 93 percent increase. If this increase is coming entirely from the extensive margin, it would represent an increase in the share withdrawing of approximately 7.5 percentage points. Alternatively, if the increase is fully due to the intensive margin, it would represent an increase in conditional annual withdrawals of approximately $19,700. As described earlier, since the distribution amounts conditional on being nonzero do not change markedly over the calendar years in our sample, it is likely that the increase we identify is coming from extensive margin responses.

5. Conclusion

Despite active research that documents pre-retirement withdrawals from retirement savings accounts, there has not been much prior work that investigates the relationship between pre-retirement withdrawal penalties and distributions from retirement accounts. One of the large barriers to understanding the effects of these penalties on distributions from retirement accounts has been data limitations, as household surveys have limited sample size and potentially underreported withdrawal activity and administrative data is often collected at longer frequencies, making it difficult to uncover event studies at shorter frequencies.

This study attempts to overcome several of these shortcomings in the data by developing new empirical techniques that allow us to analyze withdrawal activity when the penalty for early withdrawals is lifted with high-quality data from the IRS. By exploiting variation in date of birth, which leads to natural variation in exposure to penalty-free withdrawals over calendar years, we can estimate event studies that show how withdrawal behavior changes on either side of the age 59 ½ threshold.

Our results indicate large changes in withdrawal behavior as a result of crossing age 59 ½. In particular, we find that annual distributions from IRAs increase by approximately $1,600 annually, representing an increase of approximately 93 percent relative to annual withdrawals prior to age 59 ½. Our data suggest that this increase is primarily driven by additional individuals with IRA accounts accessing their funds rather than an increase in withdrawals conditional on taking distributions.

These findings suggest that the removal of the 10 percent penalty for early withdrawals at age 59 ½ does influence withdrawal behavior among individuals with IRAs. Future work will examine heterogeneity in the penalty's effect on different groups of individuals and the effect of the penalty on other financial outcomes to better understand the broader implications of policies that may change the amount or timing of the penalty.

References

Agarwal, Sumit, Jessica Pan, and Wenlan Qian. 2016. "Age of decision: Pension savings withdrawal and consumption and debt response." *SSRN 2487658*.

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 Choi, Christopher Clayton, Christopher Harris, David Laibson, and Brigitte Madrian. 2014. "Optimal Illiquidity in the Retirement Savings System." *NBER RRC Paper NB 14-05*.

Beshears, John, James J Choi, Christopher Harris, David Laibson, Brigitte C Madrian, and Jung Sakong. 2015*a*. "Self Control and Commitment: Can Decreasing the Liquidity of a Savings Account Increase Deposits?"

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

Brown, Jeffrey R, James Poterba, and David Richardson. 2014. "Do Required Minimum Distributions Matter? The Effect of the 2009 Holiday On Retirement Plan Distributions." *NBER Working Paper 20464*.

Bryant, Victoria L, Sarah Holden, and John Sabelhaus. 2010. "Qualified Retirement Plans: Analysis of Distribution and Rollover Activity." Vol. 103, 81–90, National Tax Association.

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

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.

Goda, Gopi Shah, Matthew Levy, Colleen Flaherty Manchester, Aaron Sojourner, and Joshua Tasoff. 2015. "The Role of Time Preferences and Exponential-Growth Bias in Retirement Savings." *NBER Working Paper 21482*.

Investment Company Institute. 2015. "The U.S. Retirement Market, Fourth Quarter 2014." http://www.ici.org/info/ret_14_q4_data.xls.

Mortenson, Jacob A., Heidi R. Schramm, and Andrew Whitten. 2016. "The Effect of Required Minimum Distribution Rules on Withdrawals from Traditional Individual Retirement Accounts." *Working Paper.*

Munnell, Alicia H, and Anthony Webb. 2015. "The Impact of Leakages from 401(k)s and IRAs." *Center for Retirement Research Working Paper 15-2*.

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

Sabelhaus, John. 2000. "Modeling IRA accumulation and withdrawals." *National Tax Journal*, 865–875.

About the Authors

Gopi Shah Goda is a senior fellow and deputy director of the Stanford Institute for Economic Policy Research at Stanford University. Her research focuses on the economics of aging in the United States and is designed to inform economic policymaking.

Goda's studies have examined how perceptual and behavioral biases affect retirement saving decisions and the effects of long-term care insurance on family members' work and location decisions. Before taking on her current position, she was a Robert Wood Johnson Scholar in Health Policy Research at Harvard University. She earned her B.S. in mathematics and actuarial science from the University of Nebraska – Lincoln and Ph.D. in economics from Stanford University.

Damon Jones is an assistant professor at the University of Chicago, Harris School of Public Policy. He also is a faculty research fellow at the National Bureau of Economic Research. Prior to joining Chicago Harris, he was a postdoctoral fellow at the Stanford Institute for Economic Policy Research. His work focuses on retirement savings and the labor supply of older workers, income taxation, household finance and behavioral economics. He completed his Ph.D. in economics at the University of California, Berkeley in 2009.

Shanthi Ramnath is a financial economist at the Office of Tax Analysis of the U.S. Department of the Treasury. She holds a Ph.D. in Economics from the University of Michigan. Her research broadly focuses on behavioral responses to social policies administered through the tax code.

Figure 1: Stylized pattern of distributions

Panel A: No Jump at 59 ½

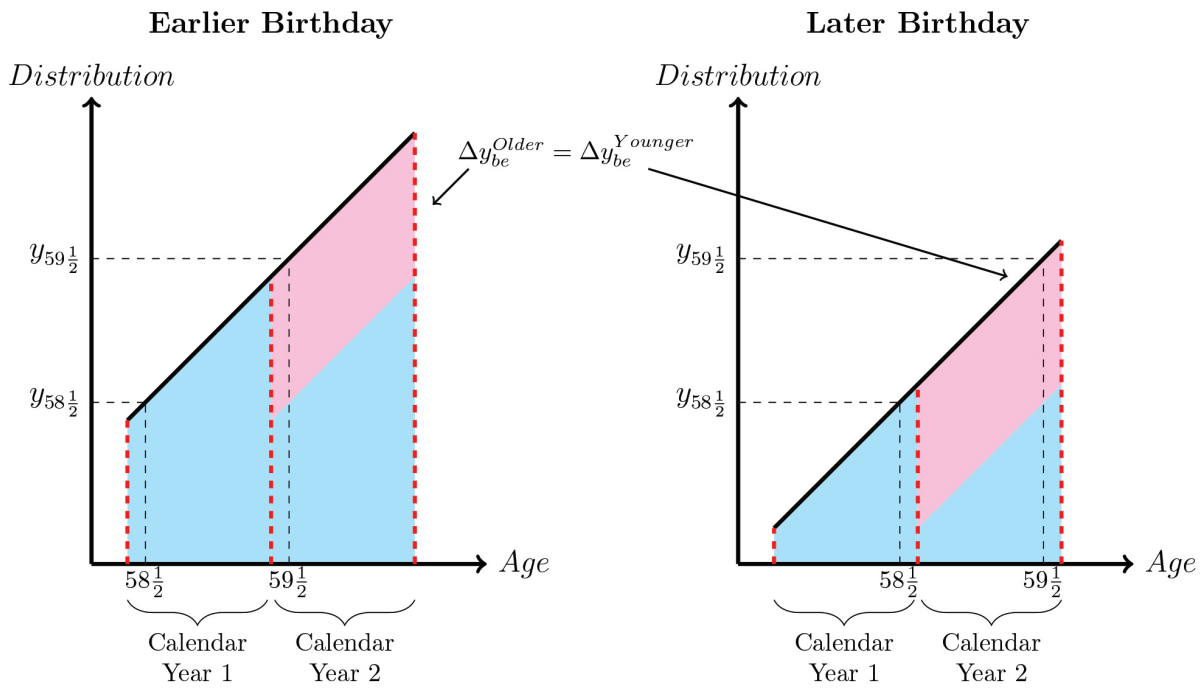

Panel B: Jump at 59 ½

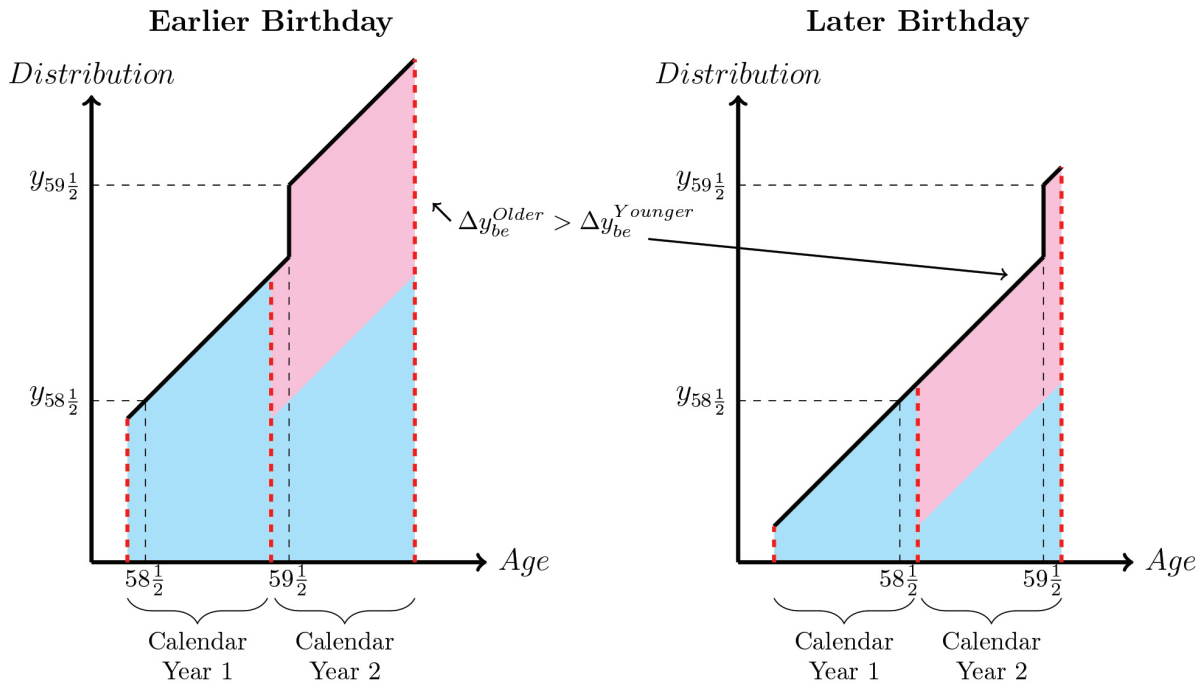

Figure 2: Mean annual IRA distribution by exposure to penalty-free withdrawal

(a) Quarterly Exposure

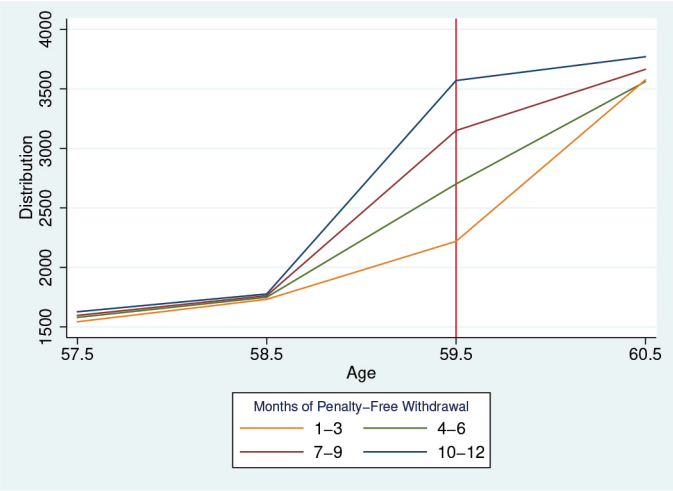

(b) Monthly Exposure

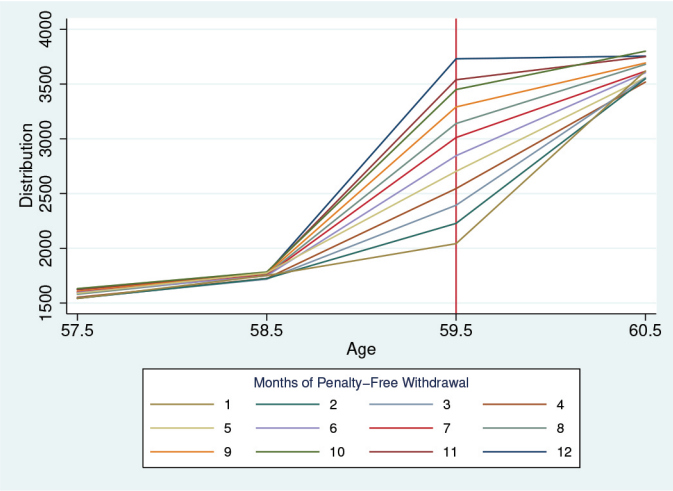

Figure 3: Average daily IRA withdrawals

(a) Linear

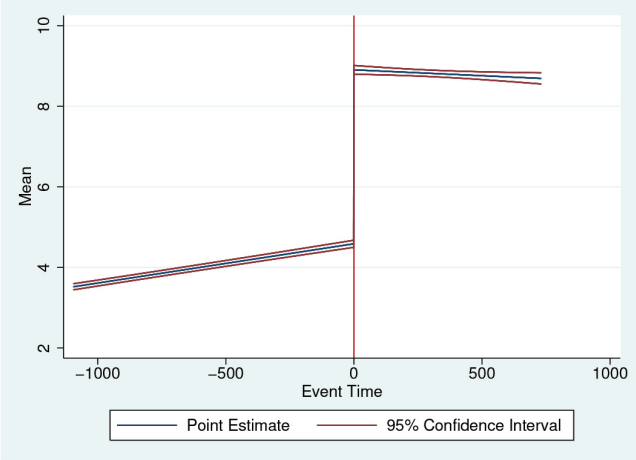

(b) Cubic

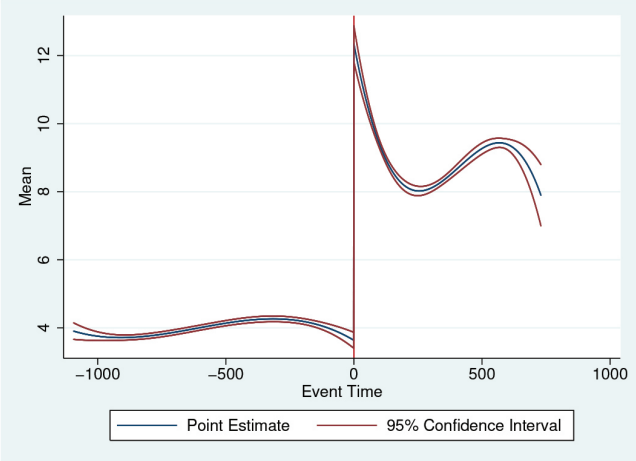

(c) Quintic

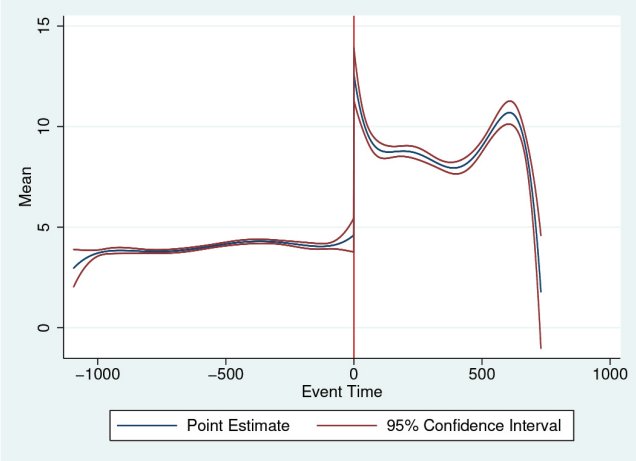

Table 1: Descriptive statistics

	57.5	58.5	59.5	60.5	All
Fraction Male	0.49	0.49	0.49	0.49	0.49
Fraction Married	0.74	0.74	0.73	0.73	0.74
Mean AGI ($)	117,855.48	116,244.20	112,832.38	108,352.78	113,821.21
Fraction with IRA Distribution	0.07	0.08	0.13	0.16	0.11
Conditional IRA Distribution Amount ($)	21,950.42	21,029.47	21,683.65	22,519.45	21,904.58
N	12,445,087	12,445,087	12,445,087	12,445,087	49,780,348

Note: Individuals born between July 1, 1941 and July 1, 1951, who have a positive fair market value of a traditional IRA account in the year they turn 57.5. Data are for the years in which an individual turns 57.5, 58.5, 59.5 and 60.5 in 1999 through 2013 tax years.

Table 2: Estimated increase in daily traditional IRA withdrawals at age 59 ½ threshold using annual patterns

	Level of Aggregation	
	Quarterly	Monthly
D_{-1}	-0.13	-0.14
	(0.02)	(0.03)
D_0	4.77	4.79
	(0.10)	(0.08)
$-D_1$	4.18	4.22
	(0.33)	(0.18)

Note: Each estimate represents the results from regressions specified according to equations 3 and 5. Under a model where daily retirement distributions increase linearly in age, we predict $\hat{D}_{-1} = 0$ and $\hat{D}_0 = -\hat{D}_1$.

Table 3: Increase in daily traditional IRA withdrawals at age 59 ½ threshold by order of polynomial and level of aggregation

	Level of Aggregation			
	Quarterly	Monthly	Weekly	Daily
Order 1	3.95	4.24	4.30	4.32
	(0.05)	(0.05)	(0.05)	(0.05)
Order 3	4.77	7.51	8.54	8.71
	(0.24)	(0.31)	(0.35)	(0.36)
Order 5	3.14	6.36	7.91	8.01
	(0.96)	(0.77)	(0.88)	(0.92)

Note: Each estimate represents the results from a separate regression with the level of aggregation given by the column header and the polynomial order given in the row. The reported estimate is the coefficient on event time 0, which represents the period in which individuals turn age 59 ½. The dependent variable in each regression is average Traditional IRA withdrawals, and the sample includes those with a positive fair market value in their IRA in the year they turn age 57 ½.

Appendix A: Simulation

Figure A.1: Simulated daily retirement distributions

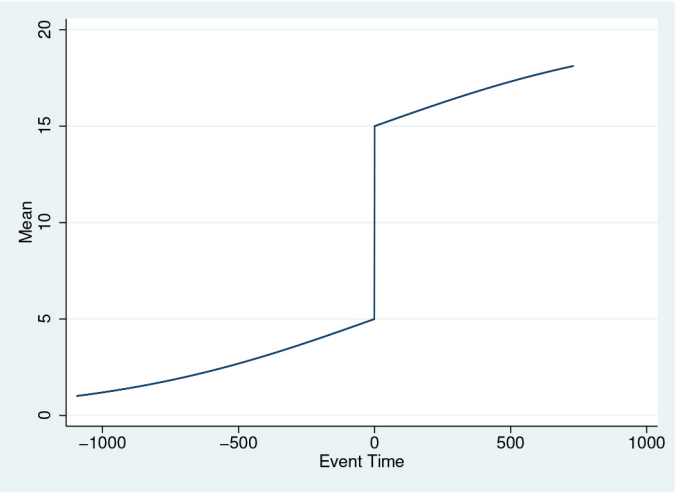

Figure A.2: Mean annual IRA distribution by exposure to penalty-free withdrawal: Simulated data

(a) Quarterly Exposure

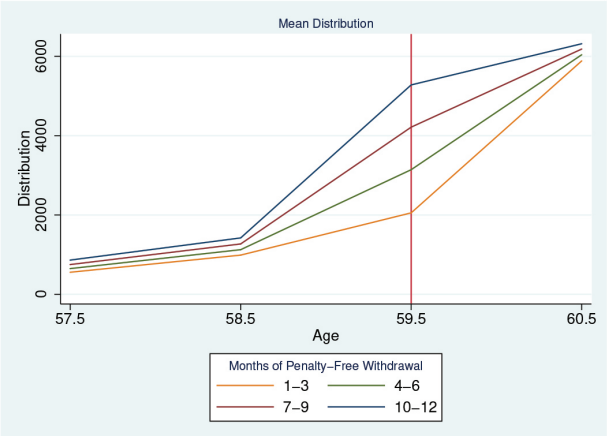

(b) Monthly Exposure

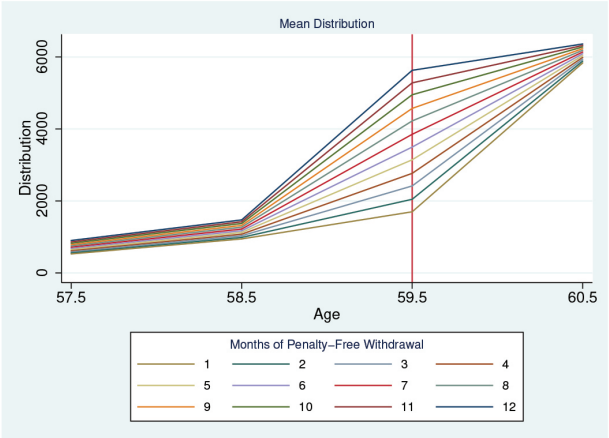

Figure A.3: Event studies using simulated annual data

(a) Linear

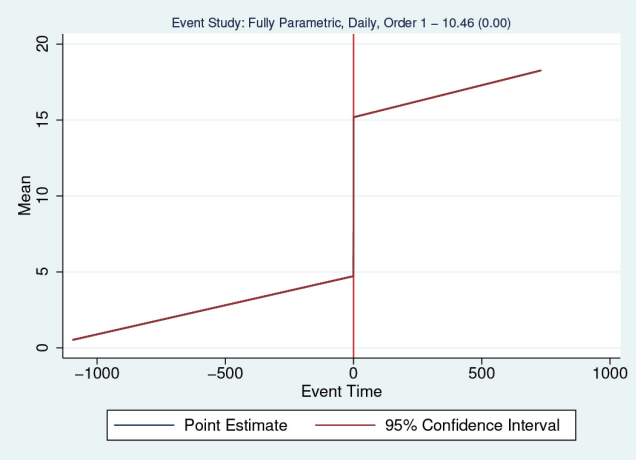

(b) Cubic

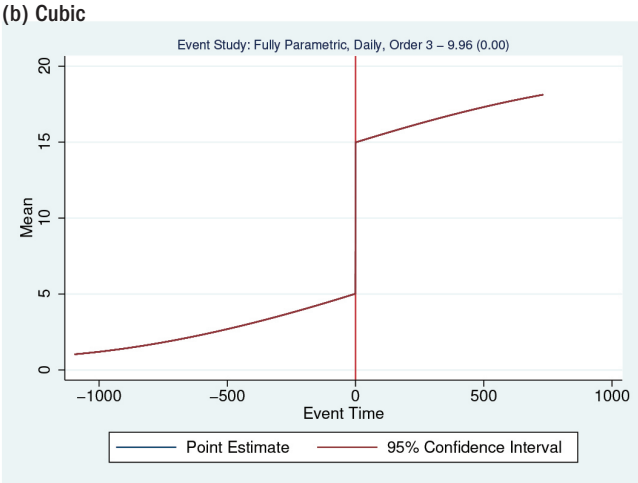

(c) Quintic

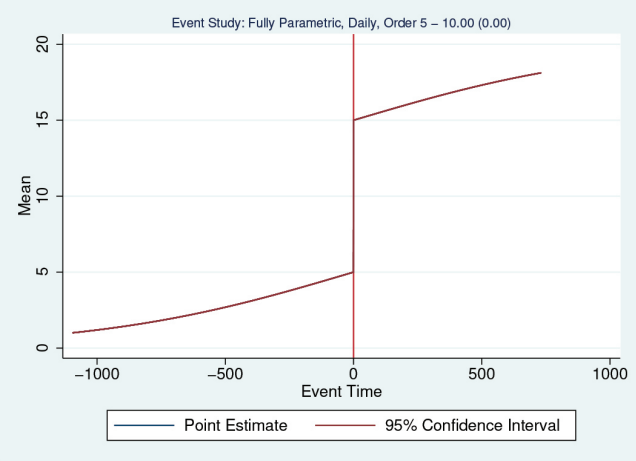

Table A.1: Estimated increase in daily traditional IRA withdrawals at age 59 ½ threshold using annual patterns: Simulated data

	Level of Aggregation	
	Quarterly	Monthly
D_{-1}	0.47	0.47
	(0.01)	(0.01
D_0	10.20	10.22
	(0.07)	(0.04)
$-D_1$	10.20	10.22
	(0.01)	(0.02)

Note: Each estimate represents the results from regressions specified according to equations 3 and 5. Under a model where daily retirement distributions increase linearly in age, we predict $\hat{D}_{-1} = 0$ and $\hat{D}_0 = -\hat{D}_1$. The data used are simulated, with a true value of $D_0 = \$10$.

Table A.2: Increase in daily withdrawals at age 59 ½ threshold by order of polynomial and level of aggregation: Simulated data

	Level of Aggregation			
	Quarterly	Monthly	Weekly	Daily
Order 1	9.29	10.17	10.40	10.46
	(0.03)	(0.01)	(0.00)	(0.00)
Order 3	5.44	8.41	9.68	9.96
	(0.10)	(0.05)	(0.01)	(0.00)
Order 5	6.05	7.31	9.70	10.00
	(0.53)	(0.17)	(0.05)	(0.00)

Note: Each estimate represents the results from a separate regression with the level of aggregation given by the column header and the polynomial order given in the row. The reported estimate is the coefficient on event time 0, which represents the period in which individuals turn age 59 ½. The dependent variable in each regression is average simulated withdrawals. The true value of the increase in daily withdrawals in the simulated data is $10.

Appendix B: Event studies for alternative frequencies

Figure B.4: Event studies for alternative frequencies: Fully parametric

(a) Weekly

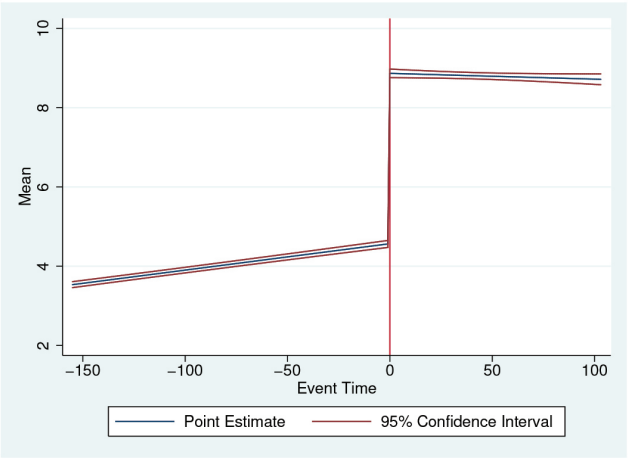

(b) Monthly

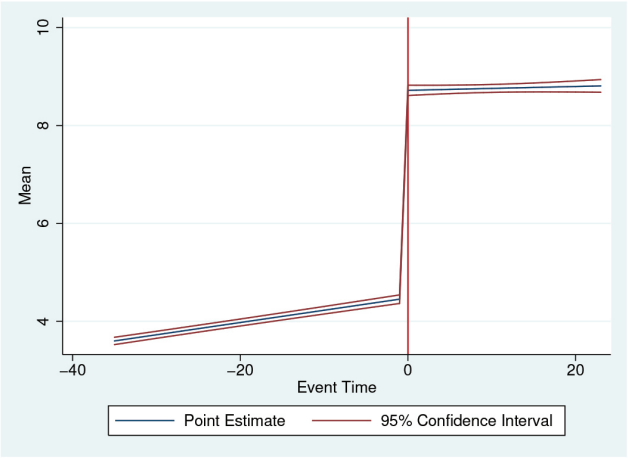

(c) Quarterly