SPRING 2018
NEW YORK UNIVERSITY
SCHOOL OF LAW

"How Do Distributions from Retirement Accounts Respond to Early Withdrawal Penalties? Evidence from Administrative Tax Returns"

Damon Jones
University of Chicago Harris School of Public Policy

March 27, 2018
Vanderbilt Hall – 208
Time: 4:00 – 5:50 p.m.
Week 9

SCHEDULE FOR 2018 NYU TAX POLICY COLLOQUIUM

(All sessions meet from 4:00-5:50 pm in Vanderbilt 208, NYU Law School)

1. Tuesday, January 16 – Greg Leiserson. Washington Center for Equitable Growth. "Removing the Free Lunch from Dynamic Scores: Reconciling the Scoring Perspective with the Optimal Tax Perspective."

2. Tuesday, January 23 – Peter Dietsch, University of Montreal Philosophy Department. "Tax Competition and Global Background Justice."

3. Tuesday, January 30 – Andrew Hayashi, University of Virginia Law School. "Countercyclical Tax Bases."

4. Tuesday, February 6 – Gerald Auten, U.S. Treasury Department. "Income Inequality in the United States: Using Tax Data to Measure Long-Term Trends."

5. Tuesday, February 13 – Vanessa Williamson, Brookings Institution. "How the Taxpaying Experience Obscures Low-Income Taxpayers and Shapes Attitudes about Progressivity"

6. Tuesday, February 27 – Jacob Goldin, Stanford Law School. "Tax Benefit Complexity and Take-up: Lessons from the Earned Income Tax Credit"

7. Tuesday, March 6 – Lisa Philipps, Osgoode Hall Law School. "Gendering the Analysis of Tax Expenditures."

8. Tuesday, March 20 – Lisa De Simone, Stanford Graduate School of Business. "Repatriation Taxes and Foreign Cash Holdings: The Impact of Anticipated Tax Reform"

9. Tuesday, March 27 – Damon Jones, University of Chicago Harris School of Public Policy. "How Do Distributions from Retirement Accounts Respond to Early Withdrawal Penalties? Evidence from Administrative Tax Returns."

10. Tuesday, April 3 – Ajay Mehrotra, American Bar Foundation and Northwestern University School of Law. "T.S. Adams and the Beginning of the Value-Added Tax."

11. Tuesday, April 10 – Jason Furman, Harvard Kennedy School. "Should Policymakers Care Whether Inequality Is Helpful or Harmful For Growth?"

12. Tuesday, April 17 – Emily Satterthwaite, University of Toronto Law School. "Electing into a Value-Added Tax: Survey Evidence from Ontario Micro-Entrepreneurs."

13. Tuesday, April 24 – Wolfgang Schon, Max Planck Institute. "Taxation and Democracy."

14. Tuesday, May 1 – Mitchell Kane, NYU Law School. "Collecting the Rent: The Global Battle to Capture MNE Profits"

How Do Distributions from Retirement Accounts Respond to Early Withdrawal Penalties? Evidence from Administrative Tax Returns*

Gopi Shah Goda
Stanford University and NBER

Damon Jones
University of Chicago and NBER

Shanthi Ramnath
U.S. Treasury

March 2018

Abstract

The design of retirement savings accounts must balance the long-term goal of retirement wealth accrual with the potential need for liquidity. Penalties (and exceptions) on pre-retirement withdrawals provide a possible lever for striking this balance. In the United States, penalties amount to 10 percent of withdrawn funds and several exceptions are available, including partial or full exemptions for the unemployed, disabled, or those incurring unreimbursed medical expenses. 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 $59\frac{1}{2}$. 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 patterns at higher levels of frequency. We find a large increase in withdrawals upon reaching age $59\frac{1}{2}$, implying an 80 percent increase in annual withdrawals on average among our population. We also show that lower-income quartiles, the recently unemployed, and those who have experienced large unreimbursed medical expenses experience larger increases, suggesting larger constraints prior to age $59\frac{1}{2}$, while those who are disabled are better able to use their account balances to smooth their consumption due to a more expansive exemption from the penalty.

JEL Classification: H24, D14, J32
Keywords: retirement savings accounts, withdrawals, distributions, penalty

*Email: gopi@stanford.edu; damonjones@uchicago.edu; shanthi.ramnath@treasury.gov. We would like to thank Andre Vasilyev and Paula Gablenz 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. In exchange for this preferential tax treatment, the accounts are relatively illiquid, as a penalty is typically imposed for withdrawals occurring before the account holder turns $59\frac{1}{2}$. This penalty is designed to dissuade people from accessing these funds prior to retirement, but there are several avenues to partially or completely liquidate funds in tax-preferred retirement savings accounts prior to retirement.[1]

The degree of illiquidity in retirement savings accounts, as determined by the rate of the penalty, the age threshold where the penalty is lifted, and the exceptions to the penalty that are granted, has implications for the accumulation of assets for retirement. First, it may affect the amount that is withdrawn from accounts prior to retirement, known as "leakage." Recent evidence suggest that leakage is 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 evidence suggests that increasing the illiquidity in retirement accounts could increase the amount of wealth accumulated for retirement by making the account a more effective commitment device.

Second, the degree of illiquidity affects the ability of individuals to smooth consump-

[1] 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 withdrawals if not paid back upon job separation. For instance, the IRS waives any penalties for workers aged 55 and older after a job termination. Finally, many accounts allow withdrawals to be taken for any reason subject to a penalty being paid.

tion by self-insuring against negative shocks. Wealth accumulated in retirement savings accounts can 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 facilitate insurance against negative income shocks, then some level of liquidity prior to retirement may be optimal. Finally, the degree of liquidity in retirement saving accounts changes the attractiveness of saving in those 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 $59\frac{1}{2}$ threshold in retirement savings accounts when the penalty for early withdrawals is removed. We assess the ability of savers to take advantage of penalty exceptions and smooth consumption in response to negative shocks by analyzing the heterogeneity in the response to the penalty across characteristics associated with shocks faced near the age threshold. If the exceptions to the penalty fully insure people from these shocks, we would expect to see higher rates of penalty-free withdrawals prior to $59\frac{1}{2}$ and smaller increases at $59\frac{1}{2}$.

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, withdrawals, 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 presents difficulties in distinguishing between general increases in retirement withdrawals as individuals age from increases occurring as a result of the removal of the penalty at age $59\frac{1}{2}$.

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 $59\frac{1}{2}$ 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 $59\frac{1}{2}$ 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 $59\frac{1}{2}$.

Our findings indicate that retirement saving illiquidity affects financial decisions. We find that increases in annual withdrawals in the calendar year one turns $59\frac{1}{2}$, relative to the previous calendar year, are larger for individuals who attain age $59\frac{1}{2}$ early in the year relative to those who attain age $59\frac{1}{2}$ late in the year. Moreover, we estimate a significant increase in pre-retirement withdrawals upon reaching age $59\frac{1}{2}$ that implies an approximately \$1,600 increase in annual withdrawals from Individual Retirement Accounts (IRAs), on average. Our results suggest that this increase is largely due to additional people taking withdrawals after the penalty is lifted rather than higher withdrawals among those who were withdrawing prior to age $59\frac{1}{2}$. This higher rate of withdrawals after $59\frac{1}{2}$ persists beyond this threshold, suggesting that the increase does not merely represent a retiming of withdrawals and that increasing liquidity by either lowering the age threshold or reducing the penalty would potentially lead to more leakage.

We also examine the heterogeneity in the response to lifting the penalty on withdrawals. Our analysis shows that individuals in higher income quartiles experience smaller increases in withdrawals at the $59\frac{1}{2}$ threshold, suggesting that the penalty creates more constraints for lower-income quartiles. We separately analyze outcomes for those who experience unemployment, disability and medical shocks and find that disability appears to be insured to a greater extent than unemployment or the risk of large unreimbursed medical expenses, as those who were recently disabled do not increase their withdrawals as much at the $59\frac{1}{2}$ threshold. This may be unsurprising given the fact that there are a broad set of exceptions to the penalty in the case of disability, while exceptions during unemployment only cover health insurance premium payments and exceptions for medical expenses only cover the amount of unreimbursed medical expenses above a threshold.

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 withdrawal amounts.

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 $70\frac{1}{2}$ 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 withdrawals for TIAA-CREF participants and find that one third of those affected by the rules discontinued their withdrawals 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 $59\frac{1}{2}$ threshold, our estimates can be interpreted as the result of the change in penalty rather than other factors. We also evaluate the role of exceptions as a form of targeted liquidity in providing consumption-smoothing benefits after negative shocks.

Second, we use a novel, comprehensive data source that provides high-quality data on withdrawals from information returns provided by the IRS. Given the relatively small numbers of individuals taking withdrawals from retirement savings accounts near the age $59\frac{1}{2}$ threshold, household surveys are unlikely to uncover any changes occurring precisely at age $59\frac{1}{2}$. In addition, household surveys may underreport withdrawals from retirement savings accounts, as even withdrawals 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 5 describes institutional features and the data we use for the study, and Section 4 lays out our empirical strategy. We discuss results in Section 6 and conclude in Section 7.

2 Background: 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 withdrawals 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 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 $59\frac{1}{2}$. 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. 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 indicate that one should 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 $59\frac{1}{2}$ on September 1 in the year in which they turn 59. In other words, the IRS considers them age $59\frac{1}{2}$ 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 $70\frac{1}{2}$ 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 $70\frac{1}{2}$.[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.

3 A Model of IRA Withdrawals

In this section, we present model designed to capture some key features of the decisionmaking environment in our setting. We begin with a simple model of IRA withdrawals and introduce a number of extensions that result in a set of stylized patterns. These basic patterns will eventually be compared to what we can estimate in the day, potentially shedding light on the key parameters of the model. We conclude with a discussion of how a more general model with richer structure might alter the predictions.

[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.

3.1 Baseline Setup

We model a finite-horizon, discrete time, life-cycle model of consumption and savings. For our purposes, it will be sufficient to restrict attention to agents' decisions in the neighborhood surrounding the time t^*, which is analogous to the date on which one reaches age $59\frac{1}{2}$. Furthermore, we abstract from earnings decisions and income shocks by assuming that agents only have an initial IRA savings account balance three years prior to time t^*, i.e. a_{t^*-3}. In this sense, we are focusing on agents who are making withdrawals from the IRA, i.e. those for whom time t^* represents a meaningful shift in incentives.

Agents have additively separable utility over consumption in each period $u\left(c_t \mid \theta_t\right)$, with $u'\left(\cdot \mid \theta_t\right) > 0$ and $u''\left(\cdot \mid \theta_t\right) < 0$. The parameter $\theta_t \in \{0, 1\}$ is an indicator for receiving a shock to marginal utility in period t. We assume $u'\left(c \mid 1\right) > u'\left(c \mid 0\right), \forall c$. For now, we also assume that the value of this "shock" is known in advance. Agents earn a gross return of R each period on savings in the IRA account and there is no borrowing. Withdrawals from the IRA are potentially penalized at the rate π_t, which takes on a value $\pi > 0$ prior to period t^* and zero thereafter. Each period, the agent solves the following:

$$
\begin{aligned}
& \max_{\{c_t\}} \sum_{j=0}^{T-t} \beta^j u\left(c_{t+j} \mid \theta_t\right) \\
\text{s.t.} \quad \left(1+\pi_t\right) c_t + a_{t+1} &= R a_t \quad \forall t \\
a_{t+1} &\geq 0 \qquad \forall t \\
a_t &> 0 \qquad \text{(given)}
\end{aligned} \tag{1}
$$

3.2 Model with no Shocks

We begin by analyzing behavior when the value of the parameter θ_t is fixed over time. The basic patterns of consumption, and by extension withdrawals, are plotted in Figure 1. In Figure 2(a) we plot consumption in the absence of the early withdrawal penalty. For simplicity, we assume that R and β are such that the agent prefers a constant level of consumption. In Figure 2(b) we introduce the early withdrawal penaly, which generates a lower level of consumption prior to period t^* relative to after. To see this, note the Euler equation between periods $t^* - 1$ and t^* is:

$$u'\left(c_{t^*-1} \mid 0\right) = (1+\pi)\,\beta R u'\left(c_{t^*} \mid 0\right) \tag{2}$$

Since we have effectively assumed that $\beta/R = 1$, then the condition implies that $c_{t^*} > c_{t^*-1}$. This generally applies to all comparisons between pre- and post-t^* consumption. Thus, the difference in levels on average before and after t^* are related to π. It follows that for agents we should expect variation in this difference based on the extent to which exceptions to the penalty are available. Finally, Figure 2(c) shows the general pattern, averaged over those who face the penalty and those who do not.

3.3 Model with Transitory Shocks

We now introduce a one-period shock to marginal utility during which $\theta_t = 1$. Relative to the prior model, consumption in a period in which a shock occurs will be *relatively* higher, both in levels and compared to neighboring time periods. In Figure 2(d) we show a shock in period $t^* - 1$ and relatively higher consumption in that same period. We can show this result by inspecting the Euler equation between period $t^* - 2$ and $t^* - 1$ with and without the shock:

$$\begin{aligned} u'\left(c_{t^*-2} \mid 0\right) &= \beta R u'\left(c_{t^*-1} \mid 0\right) \quad \text{(no shock)} \\ u'\left(c_{t^*-2} \mid 0\right) &= \beta R u'\left(c_{t^*-1} \mid 1\right) \quad \text{(shock)} \end{aligned} \tag{3}$$

Given, our assumptions, we can see that the ratio of consumption in period $t^* - 1$ to that of period $t^* - 2$ is greater when a shock occurs in period $t^* - 1$. Since marginal utility is higher, consumption must be adjusted upward in order to keep the condition satisfied. Note that π does not appear in these equations, since both periods are equally penalized. In Figure 2(e), we demonstrate a similar effect on consumption when the shock occurs in period $t^* + 1$. If we assume that the shocks are distributed i.i.d., then on average what we observe is still the difference in consumption driven by the penalty π. This is illustrated in Figure 2(f).

3.4 Model with Serial Correlated Shocks

We finally extend the model by allowing the shock to persist, now for two periods. Thus, when the shock occurs in period t it also remains in period $t+1$. In this sense, the shocks become less transitory. In Figure 2(g) we see relatively higher consumption in the periods over which the shocks occur, due to the same reasoning as before. However, we see a slightly different pattern when the shocks span period t^* in Figure 2(h). In this case, consumption is adjusted up in both periods, but there is a relatively smaller increase in consumption in period t^*-1 than in the case where the shock only lasts one period. There is also now a relatively greater increase in consumption in period t^* than in the case where the shock only lasts one period. Intuitively, resources are substituted to the two period with a shock, away from other periods, as in Figure 2(g), driven by θ. In addition, within these two periods, consumption is substituted away from period t^*-1 to period t^*, which is driven by π. Moreover, when the two period shock occurs during periods t^* and t^*+1, the increase in consumption in period t^* will be smaller, because there is no longer a wedge between the two periods created by π. It follows that if we average across individuals in this setting, we will observe the pattern in Figure 2(i). Here, we see that the serial correlation now creates a "dip" in average consumption just before period t^* and a relative "spike" in period t^*.

3.5 Discussion

4 Empirical Method

While we observe the average amount of pre-retirement withdrawals and the share taking withdrawals from IRA accounts at an annual frequency, we would like to understand how these outcomes evolve at a more granular level and, in particular, in a neighborhood near age $59\frac{1}{2}$. Our empirical strategy leverages variation in exposure to pre-retirement 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 $59\frac{1}{2}$ on December 30th, 2009, while the latter turns $59\frac{1}{2}$ on January 1st, 2010. Differences in their annual pre-retirement withdrawals and the share taking withdrawals in 2009 can be related to the fact that one person has experienced two days of penalty free withdrawals 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 withdrawal levels and withdrawal rates to test for a discontinuous effect of the age $59\frac{1}{2}$ threshold. Second, we provide a less restrictive approach that allows us to estimate event studies at age $59\frac{1}{2}$, at subannual frequencies: i.e. quarterly, monthly, weekly, and daily.

We motivate our empirical approach with a model of average daily pre-retirement withdrawals. We assume that average pre-retirement withdrawals 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 withdrawals can be characterized as follows:

$$y_{bd} = \tilde{\alpha} + \tilde{\lambda}_d + f\left(d - b - a^*\right) + D \cdot \mathbf{1}\left\{d - b \geq a^*\right\} + \varepsilon_{bd}, \tag{4}$$

where y_{bd} is the average daily withdrawal 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 $59\frac{1}{2}$, and ε_{bd} is a mean-zero error term. The function $f(\cdot)$ governs the age pattern of pre-retirement withdrawals, and its argument is measured relative to age $59\frac{1}{2}$. The function $\mathbf{1}\{\cdot\}$ is an indicator function, and, thus, the parameter D represents an additively separable shift in average pre-retirement withdrawals upon reaching age $59\frac{1}{2}$.

4.1 Annual Patterns

As a first step toward testing for a discontinuous change in behavior upon turning age $59\frac{1}{2}$, we show what can be inferred from annual patterns of pre-retirement withdrawals and the share making withdrawals. We make the extreme assumption that, aside from the possible discontinuity at age $59\frac{1}{2}$, pre-retirement withdrawals are locally linear in age, i.e. $f(j) \equiv c \cdot j$. For ease of illustration, we further assume for the moment that $\tilde{\lambda}_d = 0$.[4] Let e measure event time in years. That is, $e = 0$ in the year in which one reaches age $59\frac{1}{2}$, $e = -1$ in the year in which one turns $58\frac{1}{2}$, and so forth. We define the average annual withdrawals during event year e, for individuals 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 $59\frac{1}{2}$ is reached and event year. Within each cell, we will calculate average annual withdrawals, 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 $59\frac{1}{2}$. 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 (4), we can show the following:[5]

$$\begin{aligned} \triangle y_{q,-1} - \triangle y_{q',-1} &\approx 0 \\ \triangle y_{q,0} - \triangle y_{q',0} &\approx [q' - q] \cdot (365/4) \cdot D \\ \triangle y_{q,1} - \triangle y_{q',1} &\approx -[q' - q] \cdot (365/4) \cdot D. \end{aligned} \tag{5}$$

In words, we first measure the change in pre-retirement withdrawals from the year in which one turns $57\frac{1}{2}$ to the year in which one turns $58\frac{1}{2}$. The difference in this change across different quarters of reaching age $59\frac{1}{2}$ is approximately zero. Second, we measure the change in withdrawals from the year in which one turns $58\frac{1}{2}$ to the year in which one

[4] It is straightforward to relax this assumption by including calendar year fixed effects in the regressions below in equations (6) and (8).

[5] The expressions are only approximate for two reasons. First, the error terms do not average to zero in finite samples. Second, the differences are off by as much as 2α if any of the years involved are a leap year. The former issue vanishes asymptotically, while the latter can be handled by including calendar year fixed effects in the regressions below in quations (6) and (8).

turns $59\frac{1}{2}$. The difference in this change across different quarters of reaching age $59\frac{1}{2}$ is approximately proportional to the difference in quarters. Finally, when comparing the change in withdrawals from the year in which one turns $59\frac{1}{2}$ to the year in which one turns $60\frac{1}{2}$, the difference in average annual pre-retirement withdrawals decreases approximately linearly in the difference in quarters of reaching $59\frac{1}{2}$. We can test these predictions by estimating the following regression for different values of e:

$$\triangle y_{qe} = \alpha + D_e \cdot (-365/4)\, q + \varepsilon_{qe}, \tag{6}$$

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 $59\frac{1}{2}$ and event year, and calculate the change in average annual pre-retirement withdrawals, similarly denoted $\triangle y_{me}$, we have:

$$\begin{aligned} \triangle y_{m,-1} - \triangle y_{m',-1} &\approx 0 \\ \triangle y_{m,0} - \triangle y_{m',0} &\approx [m' - m] \cdot (365/12) \cdot D \\ \triangle y_{m,1} - \triangle y_{m',1} &\approx -[m' - m] \cdot (365/12) \cdot D. \end{aligned} \tag{7}$$

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

$$\triangle y_{me} = \alpha + D_e \cdot (-365/12)\, m + \varepsilon_{me}, \tag{8}$$

with similar predictions for $\hat{D}_e$ as in the case of quarterly averages. The above can easily be reframed to estimate the effect on the probability of making a positive withdrawal.

Figure 2 illustrates the intuition behind these results. Panel A demonstrates the pattern of pre-retirement withdrawals that would arise in the case where $D = 0$ in equation (4). The horizontal axis measures age, and the calendar years in which one turns $58\frac{1}{2}$ and $59\frac{1}{2}$ are highlighted. The vertical axis measures the average daily withdrawal. The drawing on the left depicts an “older” agent, who turns $59\frac{1}{2}$ relatively early in the year. Note, at age $59\frac{1}{2}$ there is no discrete jump in withdrawals. The area under the curve aggregates daily withdrawals into annual amounts, and the area shaded in light red measures the change in average annual withdrawals from the year in which one reaches $58\frac{1}{2}$ to the year in which one reaches $59\frac{1}{2}$, i.e. $\triangle y_{b,0}$. Similarly, the drawing on the left depicts the same patterns for a younger agent, who reaches $59\frac{1}{2}$ later in the calendar year. Although the level of annual withdrawals differs between the two, the *change* in annual withdrawals is the same.

Alternatively, Panel B of Figure 2 presents the pattern of pre-retirement withdrawals in the case where $D > 0$. Here, we see that at age $59\frac{1}{2}$, there is an upward shift in the average daily withdrawals. Furthermore, when comparing the change in withdrawals, we now see that the ”older” agent, who experiences a longer time without early withdrawal penalties, exhibits a larger change withdrawals from year to year.

Although these results rely on strong functional form assumptions, they deliver sharp predictions regarding year-to-year changes in annual withdrawals across different quarters and months of reaching $59\frac{1}{2}$. In particular, the above results imply that increases in annual withdrawals are roughly constant across quarter and/or month of reaching $59\frac{1}{2}$ between event years -2 and -1, are monotonically increasing in quarter and month of reaching $59\frac{1}{2}$ 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, allowing nonlinearity in the function $f(\cdot)$ and arbitrary values for the $\tilde{\lambda}_d$.

We can extend this approach to look at a measure of the extensive margin of withdrawals as well. We can instead use as an outcome the share of individuals in a cell making any withdrawal during the year. In this case, we can use annual patterns to learn about higher frequency patterns in the share making their *first* withdrawal of the year on a given day. Note that this is subtly different than the share of individuals making any nonzero withdrawals on a given day. Because we are limited to observing whether any withdrawals were made within a year, our method is only able to recover transitions from making no withdrawals in a year to making at least one, i.e. we cannot further identify incremental increases to say two or three withdrawals. Identifying the share of all individuals making a nonzero withdrawal on a given day would require data on the number of total number of withdrawals in a year, which we cannot observe.

4.2 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 $59\frac{1}{2}$. Formally, let $j \equiv d - b - a^*$. Let the mapping $t(b,e)$ be the calendar year in which someone born on day b reaches event year e. For example, $t(b = 10June1950, e = 0) = 2009$. 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 pre-retirement withdrawals as follows:

$$
\begin{aligned}
y_{be} &\equiv \sum_{d=\underline{d}(b,e)}^{\overline{d}(b,e)} y_{bd} \\
&= \sum_{d=\underline{d}(b,e)}^{\overline{d}(b,e)} \tilde{\alpha} + \sum_{d=\underline{d}(b,e)}^{\overline{d}(b,e)} \tilde{\lambda}_d + \sum_{j=\underline{d}(b,e)-b-a^*}^{\overline{d}(b,e)-b-a^*} f(j) + \sum_{j=\underline{d}(b,e)-b-a^*}^{\overline{d}(b,e)-b-a^*} D \cdot \mathbf{1}\{j \geq 0\} + \sum_{d=\underline{d}(b,e)}^{\overline{d}(b,e)} \varepsilon_{bd} \\
&= \alpha + \lambda_{t(b,e)} + \sum_{j=\underline{d}(b,e)-b-a^*}^{\overline{d}(b,e)-b-a^*} \left[f(j) + D \cdot \mathbf{1}\{j \geq 0\}\right] + \varepsilon_{be}, \qquad (9)
\end{aligned}
$$

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 (9). In particular, we use polynomials of order one through five, and additionally allow the coefficients to differ on either side of age $59\frac{1}{2}$. Our key parameter of interest is D, which captures any sharp change in pre-retirement withdrawals (or the share making the first withdrawal of the year) upon turning age $59\frac{1}{2}$. The method can also be adapted to model average withdrawals at lower frequencies, i.e. weekly, monthly, or quarterly.

4.3 Simulations

In Appendix A we illustrate identification with our method using simulated data of pre-retirement withdrawals. We simulate 10 cohorts of individuals, each with five years of daily withdrawals, drawn to match key moments from the actual annual retirement withdrawals in our data. Figure A.1 shows the simulated pattern of daily withdrawals two years before, and two years following age $59\frac{1}{2}$. We model a discrete jump in daily withdrawals 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 pre-retirement withdrawals by quarter and month of birth. As can be seen, the predictions in Section 4.1 are largely reflected in the simulated data. The increase in withdrawals from year to year is related the difference in exposure to penalty-free withdrawal opportunities. Next, in Figure A.3, we apply our parametric estimator of the event study. We are able to closely recreate the true, underlying pattern for daily withdrawals. In Table A.1 we report the results from the regressions in equations (6) and (8). 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 $59\frac{1}{2}$ 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 $59\frac{1}{2}$.

5 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 $57\frac{1}{2}$. This sample restriction ensures that our data contain tax years two years before and after each individual turns $59\frac{1}{2}$. Our dataset contains information on household income (Form 1040), wage earnings and employee contributions to employer-sponsored retirement plans (Form W2), withdrawals from IRAs and employer-sponsored retirement plans (Form 1099R), contributions to and account balances of IRAs (Form 5498), and tax amounts on early withdrawals (Form 5329). Because the data are unedited, 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 withdrawals 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 such as a 401(k). This makes it difficult to select a sample of individuals who are at risk of withdrawing funds from these accounts.[6] Second, while withdrawals from defined contribution plans are reported on 1099-R forms, they are undistinguishable from defined benefit payments. By contrast, IRA withdrawals 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 $59\frac{1}{2}$, while

[6] 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.

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 $59\frac{1}{2}$, 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 $57\frac{1}{2}$. 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 $57\frac{1}{2}$ in our analysis period.

Table 1 contains descriptive statistics on our sample. The data represent information from tax years in which the half-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 \$134,841. This value is relatively high both because we focus on those with assets in IRA accounts and those of older age. The fraction of our sample that takes withdrawals from their Traditional IRA is 7 or 8 percent prior to the tax year in which the individual turns age $59\frac{1}{2}$, then increases to the 16 and 17 percent in two years following the tax year when they turn $59\frac{1}{2}$. The amount withdrawn conditional on taking a distribution is approximately \$24,000 annually. The fact that this amount does not vary markedly around the age $59\frac{1}{2}$ threshold suggests that any increase in average withdrawals occurring at the age $59\frac{1}{2}$ threshold may be more likely to be on the extensive margin. Importantly, these simple comparisons of annual withdrawals across ages do not allow identification of responses to the sharp reduction in the penalty occurring when one attains age $59\frac{1}{2}$, as these increases could simply represent gradually increasing withdrawals from retirement saving accounts. The table also shows the proportion of our sample that pays a penalty on their IRA withdrawals and the average penalty amount. Approximately 5 percent of the sample incurs a penalty during the tax year in which they turn $57\frac{1}{2}$, but this declines to 1 percent at older ages.[7]

6 Results

6.1 Annual Patterns

We first investigate how annual withdrawals vary in calendar years in which individuals attain ages $57\frac{1}{2}$, $58\frac{1}{2}$, $59\frac{1}{2}$, $60\frac{1}{2}$ and $61\frac{1}{2}$ based on the exposure to penalty-free withdrawals. This exposure depends on one's quarter or month of birth, as discussed in Section 4.1. Figure 3 shows annual withdrawals from IRAs in different calendar years by exposure to penalty-free withdrawals. 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 $59\frac{1}{2}$ between August 1 and September 1 (i.e., birthdays in February).

As shown in Panel (a) of Figure 3, the level and change in annual withdrawals between the years in which age $57\frac{1}{2}$ and age $58\frac{1}{2}$ occur are the same across different quarters of birth. However, the change in annual withdrawals begins to diverge in the year individuals turn $59\frac{1}{2}$. In particular, those who have more months of penalty-free withdrawal in the calendar year in which they turn $59\frac{1}{2}$ also have larger increases in their annual withdrawals in that calendar 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 age $60\frac{1}{2}$ is reached. Accordingly, these groups now experience the greatest growth in annual withdrawals. By the time age $61\frac{1}{2}$ is reached, the gap between the groups has shrunk considerably, but not completely, suggesting that higher withdrawals for those who turned $59\frac{1}{2}$ earlier persist to some extent two years after the penalty is lifted. In Panel (b) of Figure 3, this exact result continues to hold if we group individuals by month of birth

[7]The penalty represents the additional tax as reported on the 1040 which is on a household basis, and thus may include penalties incurred by a younger spouse who has not attained age $59\frac{1}{2}$.

instead. The overall pattern is largely consistent with the predictions of Panel (b) of Figure 2.

Figure 4 shows a similar pattern when we change the outcome to the share of the sample taking any nonzero pre-retirement IRA withdrawals. The rate of withdrawals is very similar across months and quarters of penalty-free withdrawal during the tax years when people attain age $57\frac{1}{2}$ and $58\frac{1}{2}$. However, the rates change differentially based on the change in exposure to penalty free withdrawals between in the years when age $59\frac{1}{2}$ and age $60\frac{1}{2}$ are attained. Those with more exposure to penalty-free withdrawals are more likely to make a withdrawal than those with less exposure. Some of these differences persist in the years following the $59\frac{1}{2}$ tax year, but to a lesser extent than when we consider the average pre-retirement withdrawal in Figure 3.

These figures strongly suggest that the removal of the penalty at age $59\frac{1}{2}$ is driving the patterns seen in the data and largely track the predictions made in Section 4.1, where we assume a linear relationship between withdrawals and age. In particular, the increase in withdrawals when moving from age $58\frac{1}{2}$ to $59\frac{1}{2}$ is robustly monotonic in quarters or months of exposure to penalty-free withdrawals, as is predicted in the approximations in (5) and (7). To see what magnitude of increase in withdrawals at age $59\frac{1}{2}$ is implied by the figure, we estimate equations (6) and (8) in Table 2 and Table 3. Assuming a linear functional form for $f(\cdot)$, we estimate a sharp increase at age $59\frac{1}{2}$ of between \$6 and \$7 in average daily withdrawals and 0.02 basis points in the rate at which the first withdrawal of the year is made. We fail to reject the hypotheses that $D_0 = -D_1$ and $D_{-1} = 0$, both of which are predictions of the simple linear model. Overall, the patterns here are largely consistent with a discontinuous increase in withdrawals upon reaching age $59\frac{1}{2}$.

6.2 Estimated Event Study

Next we perform an event study analysis in order to trace out daily withdrawal patterns before and after individuals turn $59\frac{1}{2}$. In particular, we implement our fully-parametric approach based on equation (9), where $f(\cdot)$ is modeled using a linear, cubic, or quintic polynomial.[8] We present estimates of the daily pattern of average withdrawals in Figure 5. The vertical axes represent average daily withdrawal amounts. 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 $59\frac{1}{2}$ 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. For instance, daily average withdrawals appear to be largely flat prior to age $59\frac{1}{2}$. Thereafter, the figures show what appears to be a spike in withdrawals at age $59\frac{1}{2}$, followed by a decrease in withdrawals to a new level higher than that of prior ages. Overall, the patterns in average daily withdrawal appear to be relatively stable across choice of polynomial.[9]

Table 4 reports the results from regressions that correspond to the event study analysis using different levels of aggregation and functional form assumptions. In this table, the dependent variable captures average daily withdrawals from IRA accounts. The regression equation, in the case of a daily frequency, is shown in equation (9) and includes calendar year fixed 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 $59\frac{1}{2}$ threshold. Robust standard errors are reported below the estimated coefficient.

Using a daily level of aggregation and a linear functional form, we estimate that lifting the 10 percent penalty on withdrawals leads to an increase in the average daily withdrawal of \$4.31, which implies an annual increase of \$1,573.15. This coefficient is precisely estimated and statistically different from zero. The other reported coefficients in the table similarly

[8] In the interest of space, we omit a quadratic or quartic specification in the main text. The results are very similar when using even numbered polynomials and results from those specifications are available upon request

[9] In Appendix B, we include similar figures using different levels of aggregation — i.e. weekly, monthly, and quarterly — that show similar patterns.

show strong evidence that the removal of the penalty significantly increases withdrawals from IRAs, ranging from \$4.31 to \$9.68. As we allow for more flexibility in the polynomial, we estimate a larger jump, owing to the greater ability of the higher-order specifications to capture the spike following age $59\frac{1}{2}$.

In Figure 4 we conduct similar analysis using the annual pre-retirement withdrawal rate. Recall that in this case, our model should be interpreted as recovering the daily rate at which individuals make their first withdrawal of the year. There appears to be an increasing likelihood of making the first withdrawal as one approaches the age $59\frac{1}{2}$ threshold and then a flattening out of this pattern thereafter. The discrete increase in the likelihood of the fifth order polynomial shows very little evidence for a sharp increase at age $59\frac{1}{2}$. Overall, the patterns for the withdrawal rate are not as robustly estimated as in the case of average withdrawals. One challenge is that the annual rates of withdrawals (≈ 7 percent) translate into near zero daily rates of making the first withdrawal. Our functional form may have difficulty near this boundary and, indeed, we sometimes extrapolate to negative rates, which are theoretically impossible.

Table 5 is similar in format to Table 4, but shows regression results summarizing the estimated increase in the rate at which initial withdrawals for the year are made at age $59\frac{1}{2}$. We report the results in basis points to aid in interpretation. The estimated increase in withdrawals using the daily level of aggregation and a linear assumption suggests that approximately 2.5 out of every 10,000 individuals with IRA accounts begin withdrawing from their IRA accounts immediately after attaining age $59\frac{1}{2}$. Here, while the results are generally statistically significant, the magnitudes vary considerably across different functional form assumptions.

In Table 1, we report an unconditional average annual withdrawal amount of \$1,967.41 in the calendar year individuals turn $58\frac{1}{2}$. The implied increase of \$1,573.15 per year once one reaches age $59\frac{1}{2}$ that we estimate represents roughly an 80 percent increase from this base. The increase in the average withdrawals could arise either from the extensive margin, if a larger share of individuals are taking withdrawals, the intensive margin, if the share of individuals accessing their IRAs does not change but the average amount conditional on taking withdrawals increases, or a combination of both. In Table 5, the estimated increase in the likelihood of making the first withdrawal of the year when using a linear assumption implies a 108 percent increase in the share withdrawing relative to the average rate during the tax year when individuals turn $58\frac{1}{2}$, suggesting that the unconditional increase in average withdrawals is largely due to extensive margin responses. However, the fraction that is a result of increases in the share withdrawing is sensitive to the precise functional form assumed for the relationship between age and the probability of taking the first withdrawal.

There is some evidence from the higher order polynomials that the increase in withdrawals represents a shift in timing from periods immediately prior to the $59\frac{1}{2}$ threshold to periods immediately afterwards. However, for up to two years after this spike, a higher level of withdrawals persists. This result suggests that increasing the liquidity of retirement saving accounts by either lowering the age threshold or reducing the penalty would potentially lead to a larger amount of money withdrawn during the years that were previously covered by the early withdrawal penalty. In the next section, we investigate a different lever that changes the liquidity of retirement saving accounts, namely the exceptions provided that allow people to make withdrawals without incurring a penalty.

6.3 Heterogeneity and Penalty Exceptions

We first examine heterogeneity in the increase in average daily withdrawals at age $59\frac{1}{2}$ by quartile of adjusted gross income (AGI) and fair market value (FMV) of the IRA as reported by Form 5498 in Table 6. We report means in the tax year a person attains $58\frac{1}{2}$ of the IRA withdrawal amount (converted to a daily rate), the annual penalty paid on withdrawals, and the share of individuals who take a withdrawal for each quartile. We also report the estimated daily withdrawal amount at event time -1, just before the penalty is lifted (e.g. the day before the $59\frac{1}{2}$ birthday), and the estimated increase in daily IRA withdrawals using a cubic functional form assumption and a daily level of aggregation.

Our results suggest that both relative to the $58\frac{1}{2}$ mean and the estimated value at event time -1, the increase at $59\frac{1}{2}$ declines for those in higher income quartiles. This decline is primarily driven by smaller estimated increases at age $59\frac{1}{2}$ rather than higher levels of withdrawals under the penalty. Interestingly, those with lower levels of income are more likely to take withdrawals prior to age $59\frac{1}{2}$, but many of these early withdrawals may be exceptions to the 10 percent penalty or smaller in value, as the average penalty amount is lower for these groups. These relatively large increase in withdrawals at age $59\frac{1}{2}$ suggests that the penalty creates more constraints for those in lower-income quartiles.

The pre-withdrawal amount and rate increases dramatically by the fair market value of the accounts, as would be expected given the higher ability of larger accounts to support withdrawals. The estimated increase at $59\frac{1}{2}$ is also increasing by FMV, with the highest percentage increases occurring for the middle two quartiles. The patterns rule out the possibility that the increase in withdrawals at age $59\frac{1}{2}$ is primarily driven by the liquidation of small accounts, we observe signiifcant increases at all quartiles.

In Table 7, we examine the increases at $59\frac{1}{2}$ across different characteristics associated with negative shocks other key factors that may be related to withdrawal behavior. The tax records allow us to identify those receiving unemployment insurance, disability insurance, taking a medical deduction. We also examine differences by macroeconomic conditions by splitting our sample among those who turned $59\frac{1}{2}$ in a period of recession, where recession is defined to be between March and November 2001 or December 2007 and June 2009. In addition to these potentially negative shocks, we include a measure for having a mortgage, taking mortgage deductions, as this consumption commitment may exacerbate the effects of other negative shocks.

Many of these negative shocks trigger exceptions to the 10 percent penalty incurred on early withdrawals to different degrees. For instance, *any* early withdrawals that are made on account of disability are exempt from the penalty, while upon unemployment, only amounts paid for medical insurance for the account owner and his/her family members are exempt. Furthermore, those with unreimbursed medical expenses above 10 percent of AGI may make penalty-free IRA withdrawals to pay these expenses. There are generally no exceptions made to those who pay mortgage interest nor those who turn $59\frac{1}{2}$ in recession years.

As shown in Table 7, unemployment and high medical expenses are associated with higher-than-average withdrawals prior to $59\frac{1}{2}$ and higher-than-average penalties prior to age $59\frac{1}{2}$. By contrast, while disability is associated with slightly higher withdrawals prior to $59\frac{1}{2}$, the average penalties are actually lower than average. At the $59\frac{1}{2}$ threshold, people with recent unemployment spells or high medical expenses experience larger absolute increases, suggesting that they are constrained by the penalty. However, those who are disabled experience similar increases at $59\frac{1}{2}$ to those who are not disabled. Together, these findings suggest that the more comprehensive exemption to the penalty provided to those who are disabled allow affected individuals the ability to smooth consumption to a greater extent than those who experience unemployment or large medical expenses. We do not find strong evidence that those with mortgage interest deductions or those who turn $59\frac{1}{2}$ in recession years differ greatly in their behavior at the $59\frac{1}{2}$ threshold.

6.4 Placebo Checks

Finally, we examine whether the patterns we detect are merely mechanical artifacts of our empirical techniques by estimating placebo discontinuities at age $58\frac{1}{2}$ and $60\frac{1}{2}$ in Tables 8 and 9. The tables are in the same format as Table 4 and only include observations on either the left or the right side of the $59\frac{1}{2}$ threshold. As can be seen in these tables, only two out of the 24 estimated coefficients are statistically significant and the coefficients occasionally flip sign, while our main estimates for the $59\frac{1}{2}$ threshold show consistent evidence of an increase in withdrawals.

7 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 withdrawals from retirement accounts. One significant barrier to understanding the effects of these penalties on withdrawals 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 pre-retirement 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 $59\frac{1}{2}$ threshold.

Our results indicate large changes in withdrawal behavior as a result of crossing age $59\frac{1}{2}$. In particular, we find that the average daily withdrawal increases by approximately \$4.40, implying a \$1,600 increase in annual withdrawals from IRAs, or an increase of approximately 80 percent relative to annual withdrawals prior to age $59\frac{1}{2}$. Our data suggest that this increase is largely driven by additional individuals with IRA accounts accessing their funds rather than an increase in the average withdrawal conditional on taking withdrawals. These findings suggest that the removal of the 10 percent penalty for early withdrawals at age $59\frac{1}{2}$ does influence withdrawal behavior among individuals with IRAs.

We also investigate the heterogeneity in the response to the early withdrawal penalty and find that lower-income quartiles experience larger increases at the age $59\frac{1}{2}$ threshold in both relative and absolute terms, despite having smaller account balances on average. This finding suggests that lower-income groups within our sample are more constrained by the penalty. Finally, we show that exceptions to the penalty can allow people to insure against negative shocks, but that the ability to smooth consumption is largely affected by the scope of the penalty exemption. For instance, those who are disabled receive an exemption from penalties for any withdrawal on account of their disability, while those who are unemployed can only take withdrawals penalty-free to pay for health insurance premiums. As a result, if policymakers seek to improve wealth accumulation while still providing insurance value in the event of negative shocks, expanding the scope of penalty-free withdrawals prior to retirement may be one avenue.

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.

Figure 1: Stylized Pattern of Withdrawals

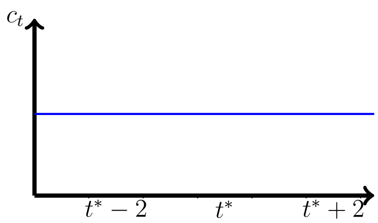

(a) No Penalty, No Shocks

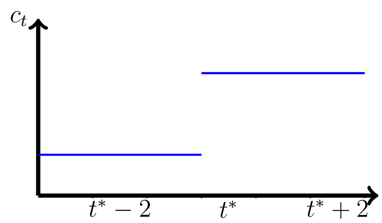

(b) Penalty, No Shocks

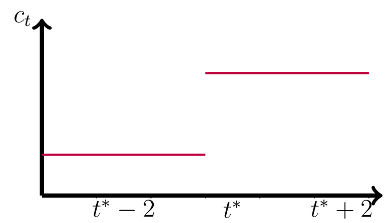

(c) Average w/ no Shocks

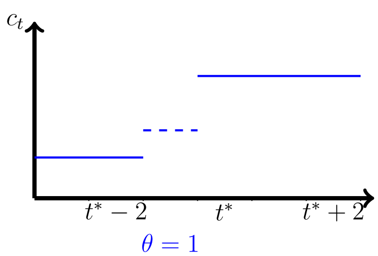

(d) Penalty w/ Shocks

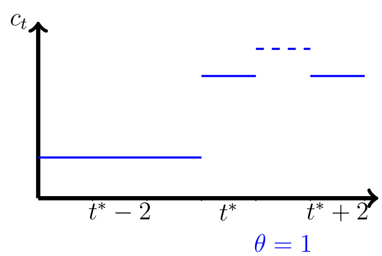

(e) Penalty w/ Shocks

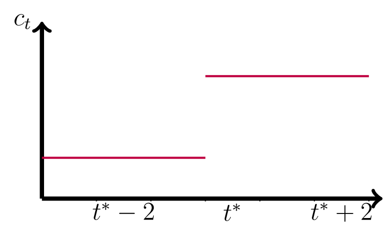

(f) Average w/ Penalty & Shocks

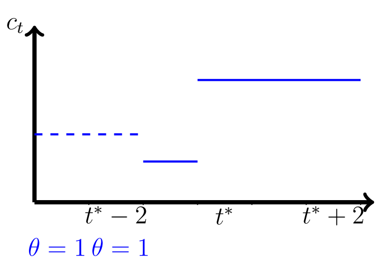

(g) Penalty & Serially Correlated Shocks

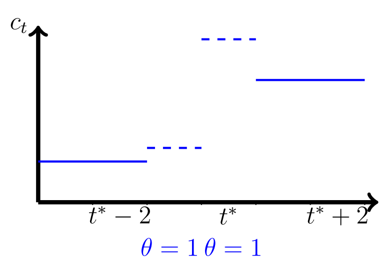

(h) Penalty & Serially Correlated Shocks

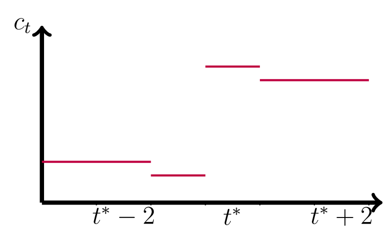

(i) Average w/ Serially Correlated Shocks

Figure 2: Stylized Pattern of Withdrawals

Panel A: No Jump at $59\frac{1}{2}$

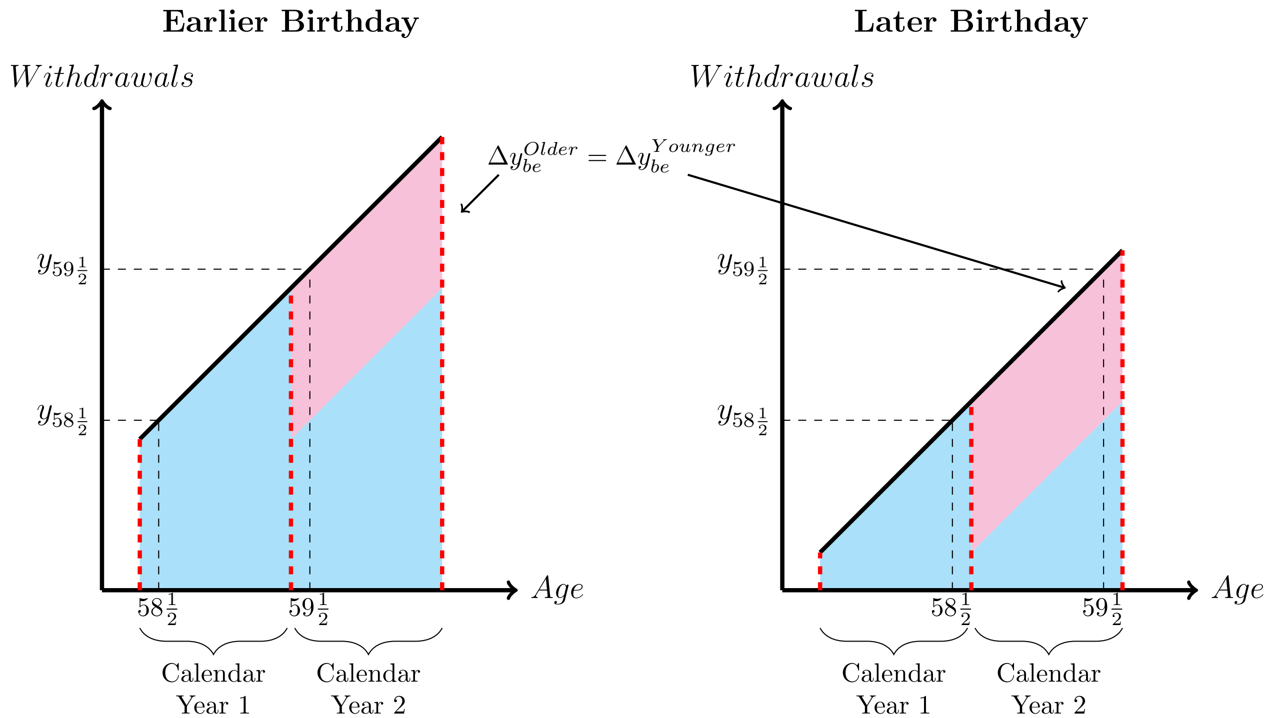

Panel B: Jump at $59\frac{1}{2}$

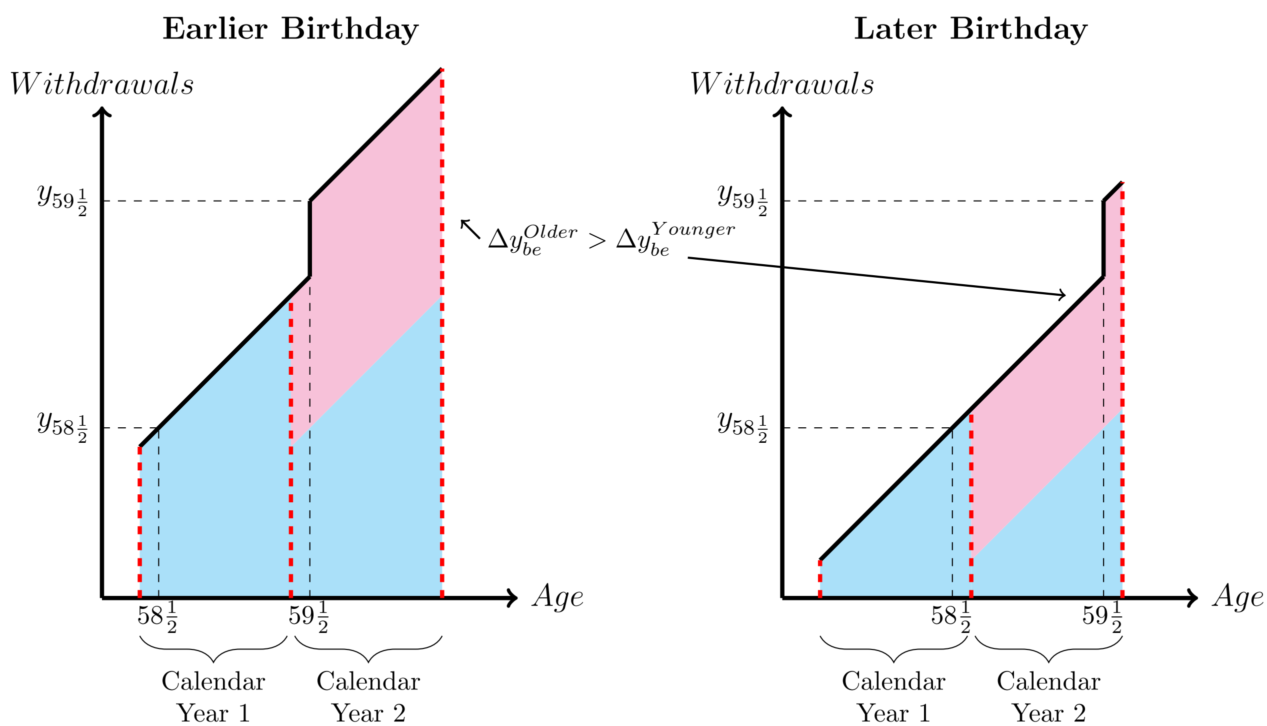

Figure 3: Mean Annual IRA Withdrawal by Exposure to Penalty-Free Withdrawal

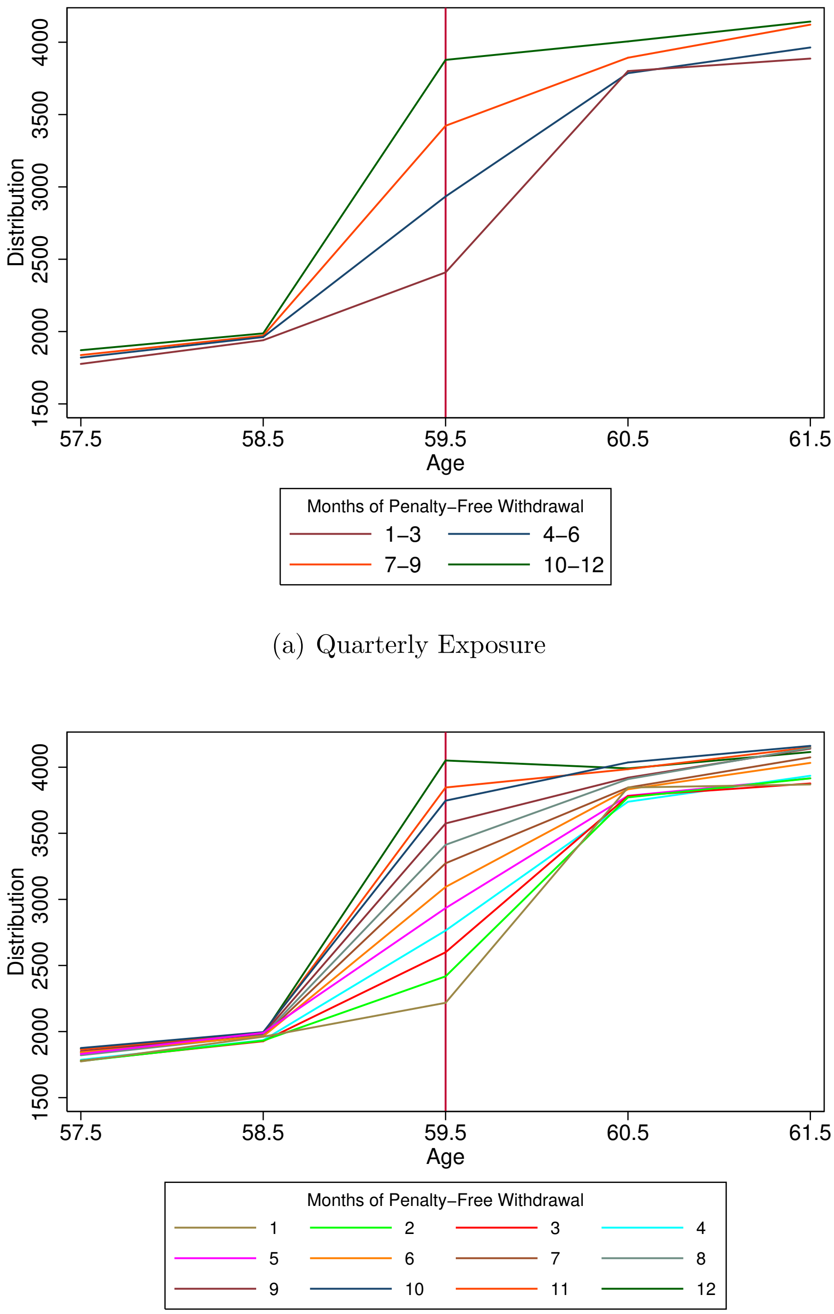

(a) Quarterly Exposure

(b) Monthly Exposure

Figure 4: Mean Annual IRA Withdrawal Rate by Exposure to Penalty-Free Withdrawal

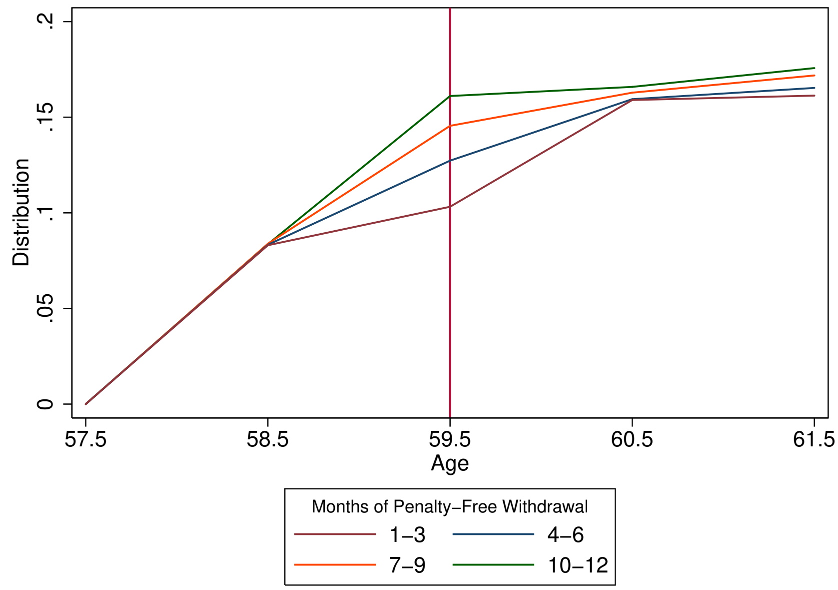

(a) Quarterly Exposure

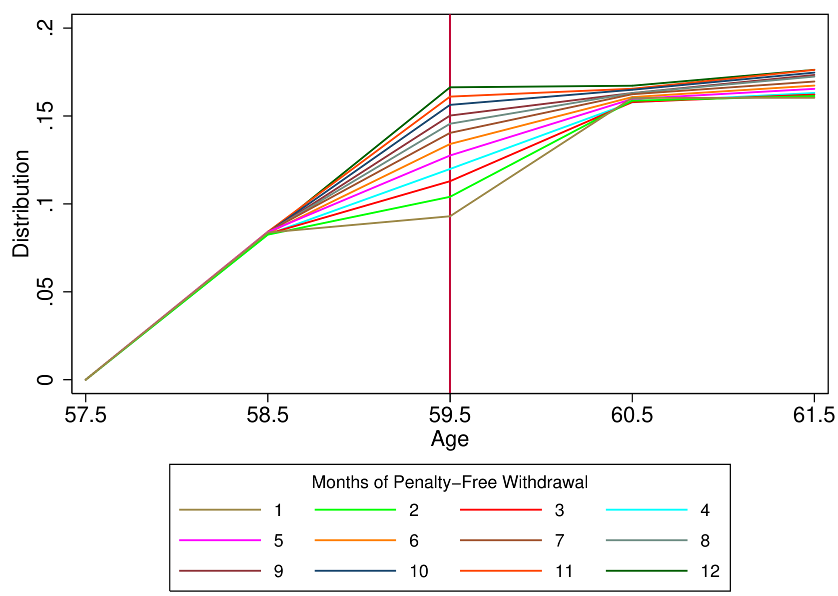

(b) Monthly Exposure

Figure 5: Average Daily IRA Withdrawals

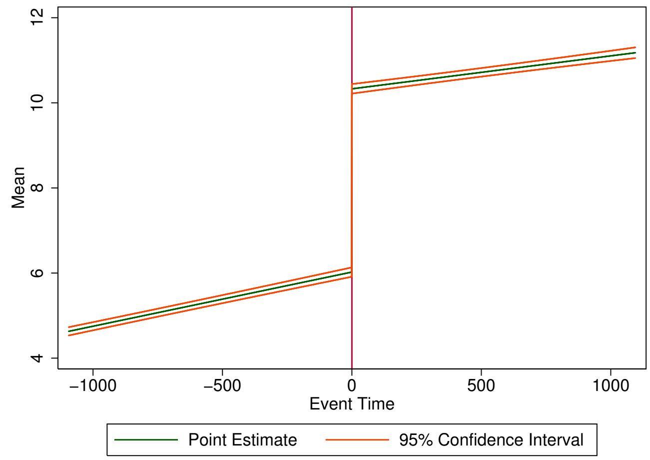

(a) Linear

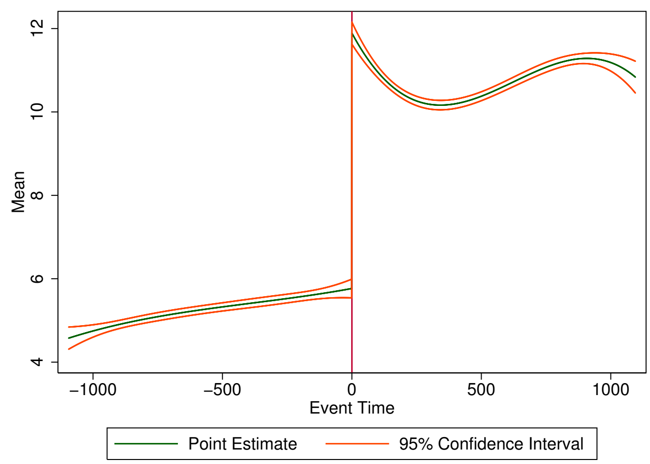

(b) Cubic

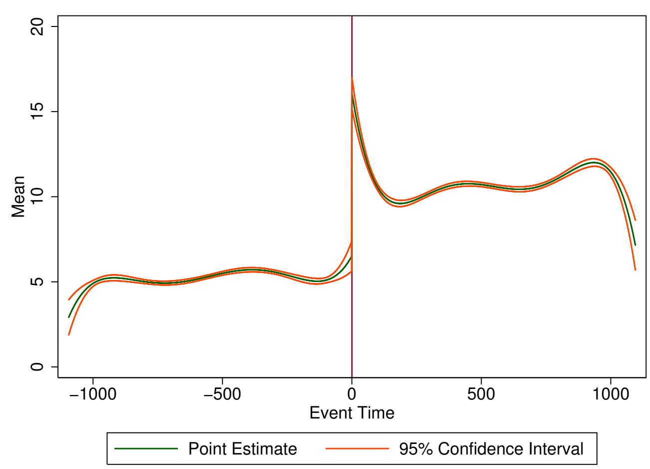

(c) Quintic

Figure 6: Average Daily IRA Withdrawal Rate (in basis points)

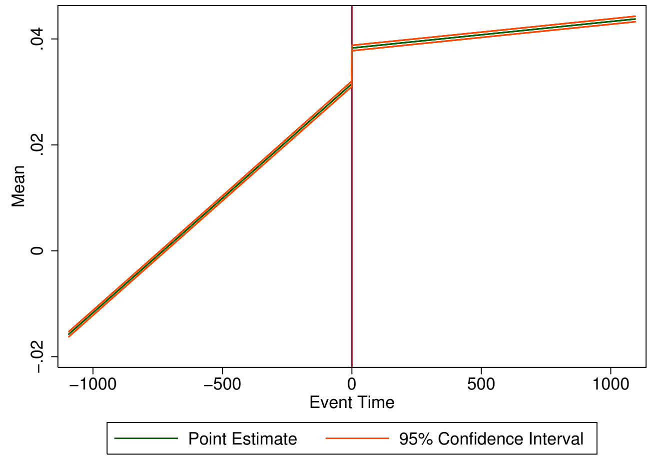

(a) Linear

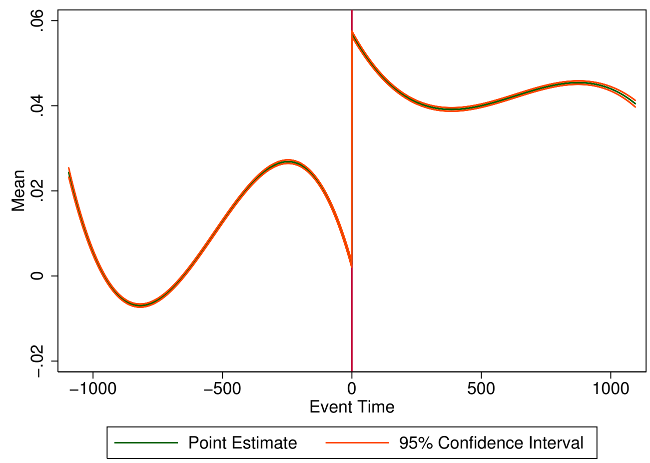

(b) Cubic

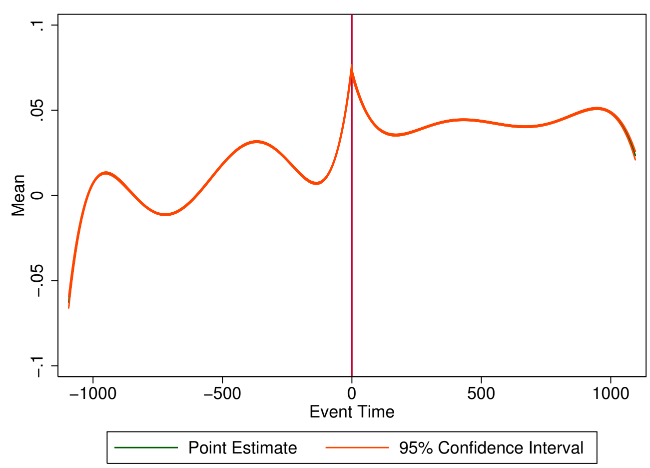

(c) Quintic

Table 1: Descriptive Statistics

	(1)	(2)	(3)	(4)	(5)
	Half-age attained during the calendar year				
	57.5	58.5	59.5	60.5	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
Traditional IRA Distribution	1827.22	1967.41	3188.56	3903.07	4081.25
Traditional Non-IRA Distribution	3892.61	4248.61	4955.11	5769.58	6207.55
Mortgage Interest (All)	5912.92	5609.80	5307.47	5324.13	4553.24
Mortgage Interest (Filers)	5048.16	4947.93	4773.65	4528.20	4259.10
Wages	47868.50	45186.34	42777.00	39808.78	37227.63
Adjusted Gross Income	134841.80	128522.28	124190.44	120452.14	117819.01
Additional Tax on IRAs	73.67	74.32	34.12	9.78	8.04
Fair Market Value of IRA	96802.69	101035.20	108727.61	112359.30	120447.29
Takes Traditional IRA Distribution	0.07	0.08	0.13	0.16	0.17
Takes Traditional Non-IRA Distribution	0.17	0.19	0.22	0.25	0.27
Has Mortgage Interest (All)	0.53	0.51	0.49	0.46	0.44
Has Mortgage Interest (Filers)	0.44	0.43	0.43	0.41	0.40
Has Wages	0.73	0.71	0.68	0.65	0.62
Has Non-Zero Adjusted Gross Income	0.96	0.95	0.94	0.93	0.92
Pays Additional Tax on IRAs	0.05	0.05	0.03	0.01	0.01
Has Non-Zero FMV on IRA	1.00	0.97	0.94	0.91	0.89
Traditional IRA Distribution (Non-Zero)	25263.76	23552.38	23555.08	23933.24	23898.01
Traditional Non-IRA Distribution (Non-Zero)	22258.54	22690.66	22683.38	22916.61	22603.59
Fair Market Value of IRA (Non-Zero)	96877.34	104078.93	115078.17	122254.66	134126.28
Number of Observations	12445149.00	12445149.00	12445149.00	12445149.00	12445149.00

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 $59\frac{1}{2}$ Threshold Using Annual Patterns

	(1)	(2)
	Level of Aggregation	
	Quarterly	Monthly
D_{-1}	-0.25	-0.25^{*}
	(0.22)	(0.12)
D_0	6.95^{***}	6.99^{***}
	(0.09)	(0.20)
$-D_1$	6.18^{***}	6.31^{***}
	(0.19)	(0.29)

Note: Each estimate represents the results from regressions specified according to equations (6) and (8). Under a model where daily retirement withdrawals increase linearly in age, we predict $\hat{D}_{-1} = 0$ and $\hat{D}_0 = -\hat{D}_1$.

Table 3: Estimated Increase in Daily Traditional IRA Withdrawal Rate (in basis points) at Age $59\frac{1}{2}$ Threshold Using Annual Patterns

	(1)	(2)
	Level of Aggregation	
	Quarterly	Monthly
D_{-1}	0.0002	0.0002
	(0.0001)	(0.0002)
D_0	0.0208***	0.0209***
	(0.0014)	(0.0010)
$-D_1$	0.0184*	0.0185***
	(0.0019)	(0.0015)

Note: Each estimate represents the results from regressions specified according to equations (6) and (8). Under a model where daily retirement withdrawals increase linearly in age, we predict $\hat{D}_{-1} = 0$ and $\hat{D}_0 = -\hat{D}_1$.

Table 4: Increase in Daily Traditional IRA Withdrawals at Age $59\frac{1}{2}$ Threshold by Order of Polynomial and Level of Aggregation

	(1)	(2)	(3)	(4)
	Time Period			
	Quarterly	Monthly	Weekly	Daily
Order 1	4.36	4.37	4.32	4.31
	(0.04)	(0.04)	(0.04)	(0.04)
Order 3	5.29	6.19	6.14	6.13
	(0.13)	(0.18)	(0.20)	(0.21)
Order 5	4.81	8.48	9.68	9.60
	(0.26)	(0.58)	(0.76)	(0.81)

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 $59\frac{1}{2}$. 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 $57\frac{1}{2}$.

Table 5: Increase in Daily Traditional IRA Withdrawal Rate (in basis points) at Age $59\frac{1}{2}$ Threshold by Order of Polynomial and Level of Aggregation

	(1)	(2)	(3)	(4)
	Time Period			
	Quarterly	Monthly	Weekly	Daily
Order 1	0.0358	0.0289	0.0257	0.0248
	(0.0004)	(0.0004)	(0.0004)	(0.0004)
Order 3	0.1379	0.1824	0.1941	0.1977
	(0.0012)	(0.0016)	(0.0019)	(0.0019)
Order 5	0.0801	0.0524	0.0067	-0.0097
	(0.0018)	(0.0040)	(0.0055)	(0.0058)

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 $59\frac{1}{2}$. The dependent variable in each regression is the proportion of each date of birth cell who makes Traditional IRA withdrawals, and the sample includes those with a positive fair market value in their IRA in the year they turn age $57\frac{1}{2}$.

Table 6: Heterogeneity in Increase in Average Withdrawals at $59\frac{1}{2}$ by Income and Account Value

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Means at 58.5					Percent Increase	
	IRA Distribution (Daily)	Penalty (Annual)	Takes IRA Withdrawal	Estimated Mean at T-1	Increase at 59.5	Relative to 58.5 Mean	Relative to T-1 Mean
Full Sample	5.39	74.32	0.08	5.77	6.13	113.7%	106.2%
				(0.11)	(0.21)		
By Adjusted Gross Income							
First Quartile	5.05	67.90	0.12	4.44	7.6	150.6%	171.2%
				(0.16)	(0.30)		
Second Quartile	5.18	68.64	0.09	5.27	6.04	116.6%	114.6%
				(0.14)	(0.26)		
Third Quartile	5.51	74.07	0.08	6.05	5.72	103.9%	94.5%
				(0.18)	(0.31)		
Fourth Quartile	5.83	86.66	0.05	8.35	4.97	85.3%	59.5%
				(0.36)	(0.64)		
By Fair Market Value							
First Quartile	0.78	38.56	0.06	1.52	0.68	86.9%	44.7%
				(0.08)	(0.13)		
Second Quartile	1.67	46.50	0.06	1.37	3.93	235.4%	286.9%
				(0.08)	(0.17)		
Third Quartile	3.39	67.08	0.07	2.32	7.18	211.9%	309.5%
				(0.13)	(0.24)		
Fourth Quartile	15.71	145.11	0.15	18.82	14.24	90.6%	75.7%
				(0.46)	(0.84)		

Table 7: Heterogeneity in Increase at $59\frac{1}{2}$ by Negative Shocks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
		Means at 58.5					Percent Increase	
	Share	IRA Distribution (Daily)	Penalty (Annual)	Takes IRA Withdrawal	Estimated Mean at T-1	Increase at 59.5	Relative to 58.5 Mean	Relative to T-1 Mean
Full Sample		5.39	74.32	0.08	5.77	6.13	113.7%	106.2%
					(0.11)	(0.21)		
By Receipt of UI at 57.5								
Received UI at 57.5	0.05	8.04	139.38	0.15	5.55	11.93	148.3%	215.0%
					(0.45)	(0.77)		
Did not receive UI at 57.5	0.95	5.26	71.14	0.08	5.79	5.84	111.0%	100.9%
					(0.12)	(0.22)		
By Receipt of DI at 57.5								
Received DI at 57.5	0.03	6.56	47.63	0.14	6.17	5.78	88.1%	93.7%
					(0.83)	(0.96)		
Did not receive DI at 57.5	0.97	5.35	75.26	0.08	5.75	6.14	114.8%	106.8%
					(0.12)	(0.21)		
By Medical Deduction at 57.5								
Took deduction	0.07	9.56	121.48	0.13	9.26	8.12	84.9%	87.7%
					(0.48)	(0.82)		
Did not take deduction	0.93	5.07	70.66	0.08	5.55	6.01	118.6%	108.3%
					(0.12)	(0.21)		
By Mortgage Interest at 57.5								
Took deduction	0.63	6.22	92.79	0.09	6.96	6.9	110.9%	99.1%
					(0.16)	(0.28)		
Did not take deduction	0.37	3.98	43.00	0.08	4.14	4.65	116.7%	112.3%
					(0.15)	(0.29)		
By Turns 59.5 in Recession Year								
Non-recession year	0.76	5.37	72.45	0.08	5.74	6.02	112.2%	104.9%
					(0.15)	(0.24)		
Recession year	0.24	5.46	80.39	0.08	5.84	6.65	121.7%	113.9%
					(0.27)	(0.5)		

Table 8: Placebo: Increase in Daily Traditional IRA Withdrawals at Age $58\frac{1}{2}$ Threshold by Order of Polynomial and Level of Aggregation

	Time Period			
	Quarterly	Monthly	Weekly	Daily
Order 1	0.34	0.26	0.23	0.22
	(0.10)	(0.13)	(0.14)	(0.14)
Order 3	0.33	0.78	0.29	0.21
	(0.29)	(0.63)	(0.81)	(0.87)
Order 5	0.28	1.75	1.13	0.56
	(0.32)	(1.01)	(2.16)	(2.59)

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 $58\frac{1}{2}$. 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 $57\frac{1}{2}$.

Table 9: Placebo: Increase in Daily Traditional IRA Withdrawals at Age $60\frac{1}{2}$ Threshold by Order of Polynomial and Level of Aggregation

	Time Period			
	Quarterly	Monthly	Weekly	Daily
Order 1	0.36	-0.18	-0.40	-0.47
	(0.18)	(0.20)	(0.22)	(0.22)
Order 3	3.02	1.32	-0.06	-0.68
	(0.42)	(0.95)	(1.36)	(1.46)
Order 5	3.02	1.14	-0.63	-1.81
	(0.52)	(1.47)	(3.54)	(4.25)

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 $60\frac{1}{2}$. 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 $57\frac{1}{2}$.

Appendix A: Simulation

Figure A.1: Simulated Daily Retirement Withdrawals

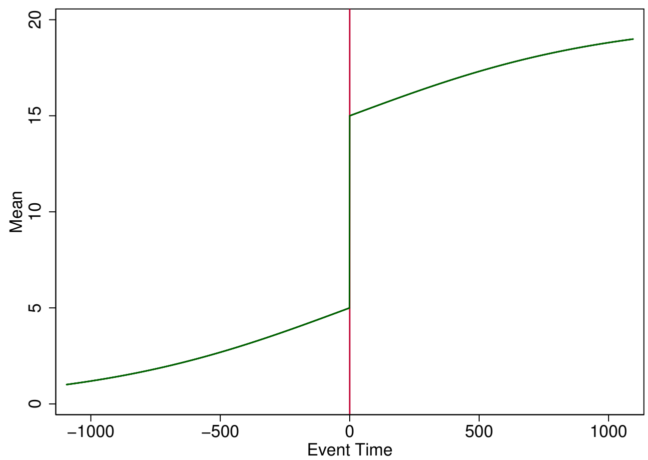

Figure A.2: Mean Annual IRA Withdrawal by Exposure to Penalty-Free Withdrawal: Simulated Data

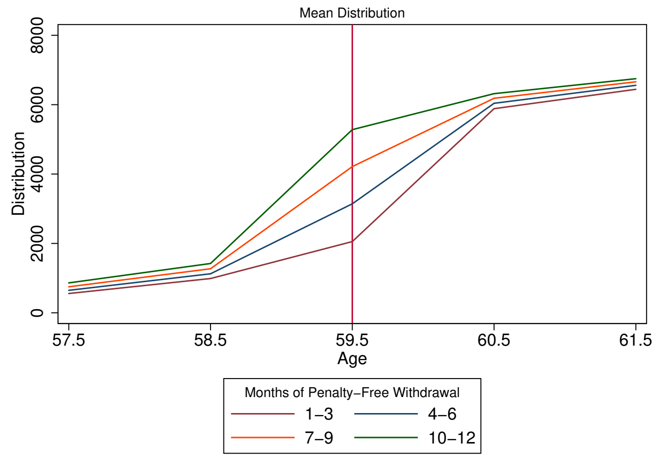

(a) Quarterly Exposure

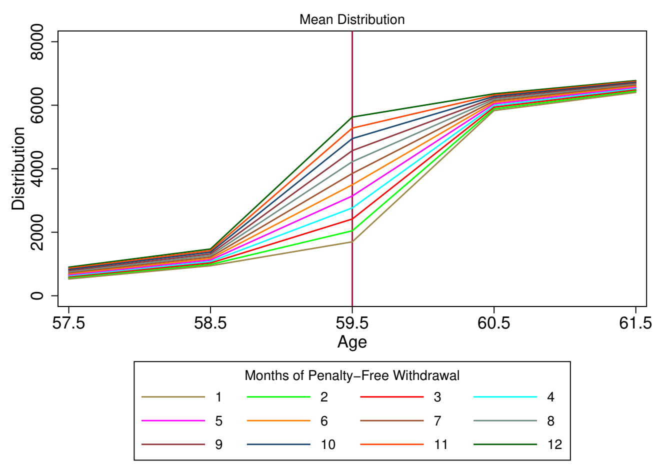

(b) Monthly Exposure

Figure A.3: Event Studies Using Simulated Annual Data

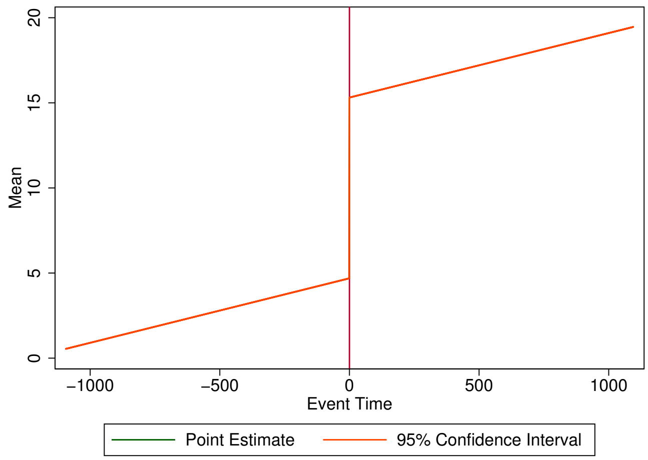

(a) Linear

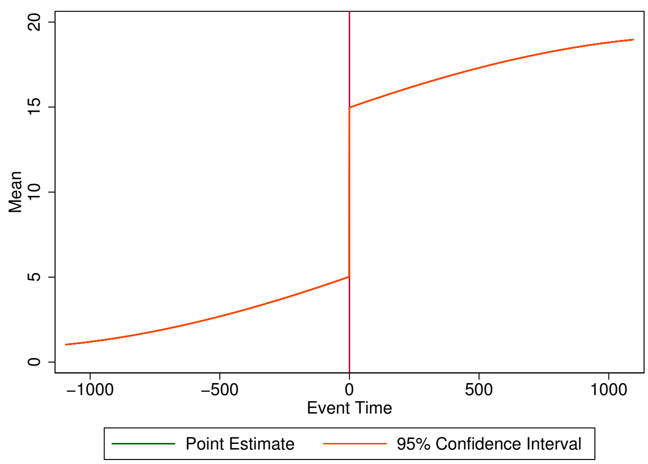

(b) Cubic

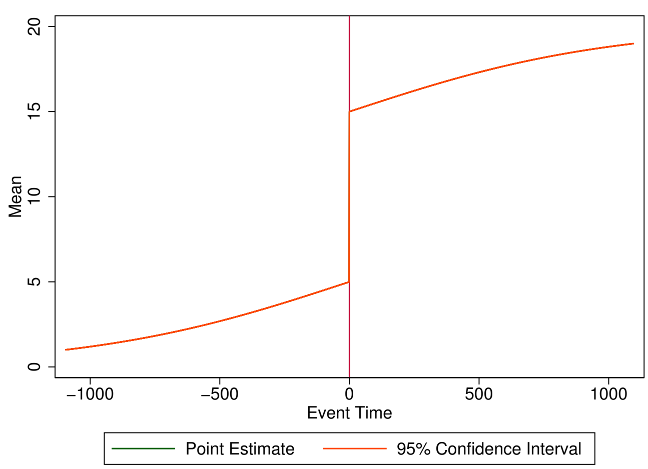

(c) Quintic

Table A.1: Estimated Increase in Daily Traditional IRA Withdrawals at Age $59\frac{1}{2}$ Threshold Using Annual Patterns

	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)
$-D_1$	10.2***	10.22***
	(0.01)	(0.02)

Note: Each estimate represents the results from regressions specified according to equations 6 and 8. 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 A.2: Increase in Daily Traditional IRA Withdrawals at Age $59\frac{1}{2}$ Threshold by Order of Polynomial and Level of Aggregation: Simulated Data

	Time Period			
	Quarterly	Monthly	Weekly	Daily
Order 1	10.22	10.58	10.61	10.62
	(0.03)	(0.01)	(0.00)	(0.00)
Order 3	8.15	9.69	9.88	9.95
	(0.10)	(0.04)	(0.01)	(0.00)
Order 5	4.21	8.03	9.90	10.00
	(0.57)	(0.15)	(0.04)	(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 $59\frac{1}{2}$. 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 $57\frac{1}{2}$.

Appendix B: Event Studies for Alternative Frequencies

Figure B.1: Event Studies for Alternative Frequencies and Functional Forms: Daily IRA Withdrawals

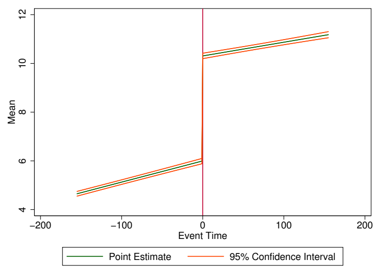

(a) Weekly, Linear

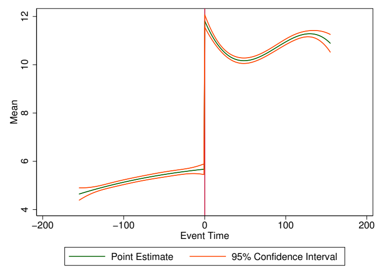

(b) Weekly, Cubic

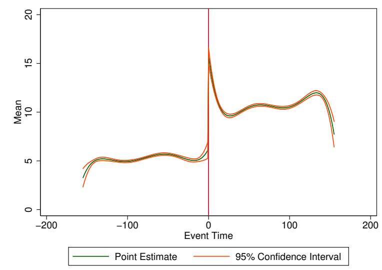

(c) Weekly, Quintic

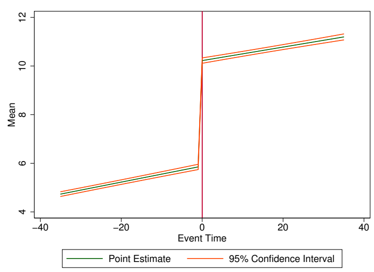

(d) Monthly, Linear

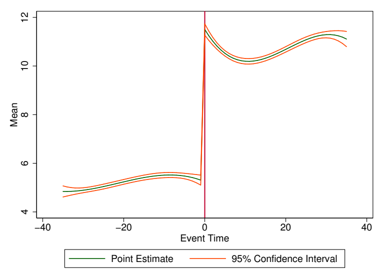

(e) Monthly, Cubic

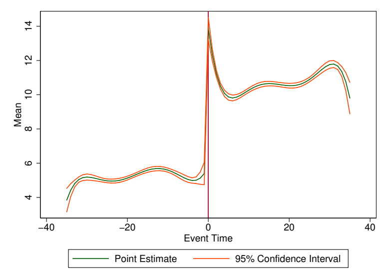

(f) Monthly, Quintic

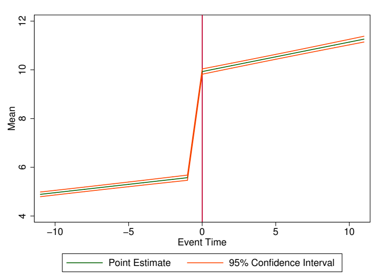

(g) Quarterly, Linear

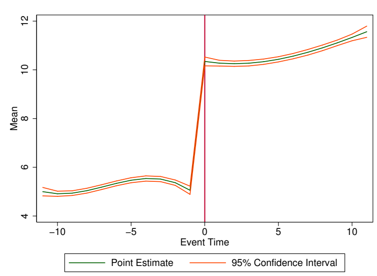

(h) Quarterly, Cubic

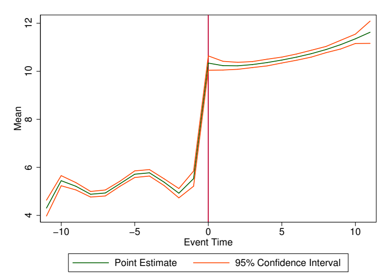

(i) Quarterly, Quintic

Figure B.2: Event Studies for Alternative Frequencies and Functional Forms: Daily IRA Withdrawal Rate

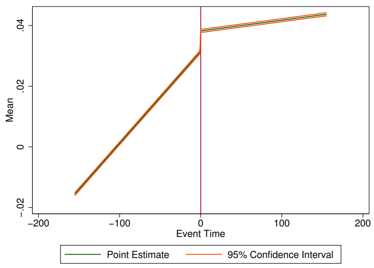

(a) Weekly, Linear

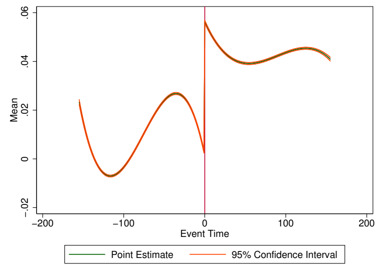

(b) Weekly, Cubic

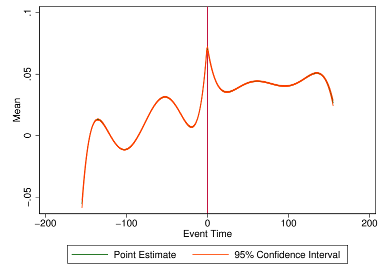

(c) Weekly, Quintic

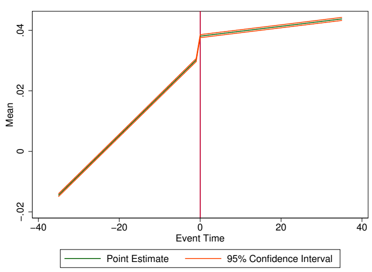

(d) Monthly, Linear

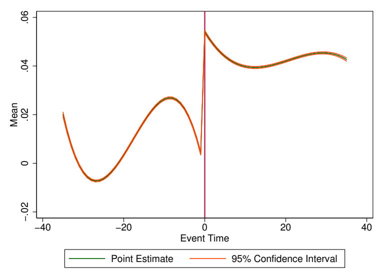

(e) Monthly, Cubic

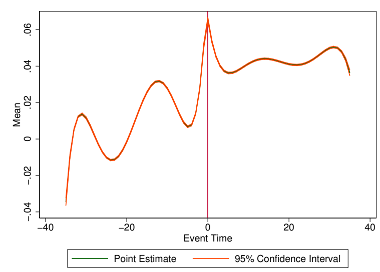

(f) Monthly, Quintic

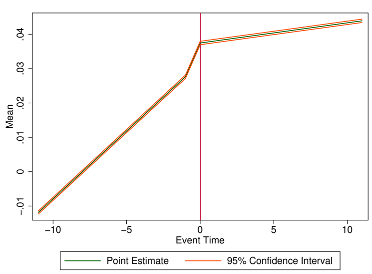

(g) Quarterly, Linear

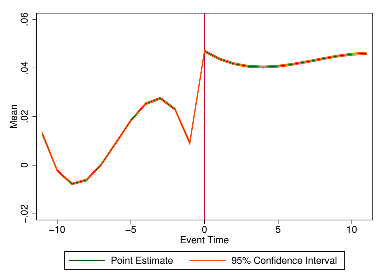

(h) Quarterly, Cubic

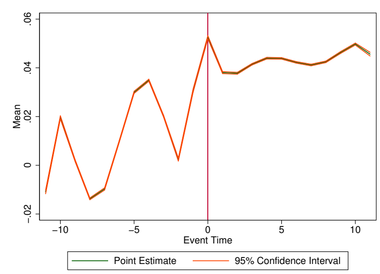

(i) Quarterly, Quintic