The Impact of Withdrawal Penalties on Retirement Savings

Ellen Stuart
Harvard University
Job Market Paper

Victoria Bryant
Internal Revenue Service
Statistics of Income Division

This version: November 16, 2021
Latest version available here

Tax-benefited retirement accounts have features designed to encourage saving, including a penalty for withdrawing before age $59\frac{1}{2}$. Account holders also face a penalty for failing to take required minimum withdrawals after age 72. Using a bunching analysis, we estimate that these penalties cause over 17% of traditional IRA holders to change their withdrawal timing each year, shifting almost $60 billion of distributions annually. We estimate a dynamic life-cycle model to analyze the effect of changing these penalties. For both penalties, we find alternative combinations of age threshold and penalty rate that lead to increased average welfare and lifetime tax remittances.

JEL classification: D15, H2, J32
Keywords: retirement savings, retirement income, tax penalties

---

Stuart: estuart@fas.harvard.edu (corresponding author). Bryant: Victoria.L.Bryant@irs.gov. We are indebted to the generous and unwavering support of Joel Slemrod, Jim Hines, Toni Whited, Charlie Brown, and Ashley Craig. We thank Mattan Alalouf, Will Boning, Giacomo Brusco, Nathaniel Hendren, Yeliz Kaçamak, Stephanie Karol, Jonathan Leganza, David Love, Jennifer Mayo, Ben Meiselman, Jake Mortenson, Dhiren Patki, Mike Ricks, and Eleanor Wilking, as well as seminar participants at the University of Michigan, the NTA Annual Conference, and the Annual Congress of the International Institute of Public Finance, for helpful comments and discussion. We're grateful to our colleagues at the IRS, particularly John Guyton and Alicia Miller. Aaron Bauer, Jonathan Kummerfeld, and Dhiren Patki provided invaluable technical support. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the Internal Revenue Service. All results have been reviewed by the IRS to ensure that no confidential information is disclosed. All errors are our own.

Stuart gratefully acknowledges financial support from the NBER Pre-Doctoral Fellowship in Retirement and Disability Policy Research (under Director Nicole Maestas). The research reported herein was performed pursuant to grant RDR18000003 from the U.S. Social Security Administration (SSA) funded as part of the Retirement and Disability Research Consortium. The opinions and conclusions expressed are solely those of the authors and do not represent the opinions or policy of SSA, any agency of the Federal Government, or NBER. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of the contents of this report. Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation or favoring by the United States Government or any agency thereof.

There is growing concern about the state of retirement security in the United States. In 2015, the Government Accountability Office found that a third of U.S. households with the head of house aged 55 or older had no retirement income other than Social Security, and a quarter had no retirement income other than Social Security and a pension (Government Accountability Office 2015). For households with additional retirement savings, the median account balance translated into a monthly annuity of just $400.

One way the U.S. government encourages individuals to improve their retirement finances is by offering tax benefits to certain types of retirement savings accounts, including Individual Retirement Accounts (IRAs). IRAs are subject to two tax penalties. To discourage premature spending out of these accounts, withdrawals before age $59\frac{1}{2}$ are penalized with an additional 10% tax (the "early withdrawal penalty"). To minimize the tax cost of IRAs, account holders are required to make annual minimum withdrawals beginning at age $70\frac{1}{2}$.[1] Failure to take these required withdrawals results in an additional 50% tax on the amount not withdrawn (the "excess accumulation penalty"). Although these penalties have existed for decades, their effect on retirement savings, welfare, and tax remittances is not well understood.

This paper makes two contributions. First, we use reduced-form bunching methods to estimate how many people change the timing of withdrawals from traditional IRAs in response to the two withdrawal penalties. Our analysis suggests that, each year, 1.4% of traditional IRA holders are taking a distribution later than they would have as a result of the early withdrawal penalty, while 16.2% of traditional IRA holders take withdrawals when they otherwise would not have due to to the excess accumulation penalty. These responses shift the timing of nearly $60 billion worth of withdrawals from traditional IRAs each year.

Second, we estimate a dynamic life-cycle model, which provides a framework for analyzing the effect of changing the early withdrawal and excess accumulation penalties on welfare and tax remittances. The estimation routine yields estimates of four preference parameters: the elasticity of intertemporal substitution (EIS), the discount factor, and two bequest motive parameters. Identification of the preference parameters is driven by exogenous budget set discontinuities generated by the two tax penalties.

For both penalties, we identify combinations of the age threshold and penalty rate that lead to increased average welfare and tax remittances relative to the base policy. For the early withdrawal penalty, these combinations involve increasing the age threshold from $59\frac{1}{2}$ to $62\frac{1}{2}$ through $65\frac{1}{2}$, and lowering the penalty rate from 10% to 5%. These alternative

[1] In 2019, the age for annual minimum withdrawals was raised to 72.

combinations also yield higher IRA balances at age $65\frac{1}{2}$, furthering the purported aim of the early withdrawal penalty: to increase retirement savings. The intuition for this finding is as follows: conditional on having these penalties in place, there are benefits to encouraging taxpayers to keep their money in these accounts as long as possible. However, if individuals need to take early distributions as the result of an unexpected income shock, they can do so with minimal sanction.

For the excess accumulation penalty, we find that increasing the age for required withdrawals from $70\frac{1}{2}$ to $71\frac{1}{2}$ through $73\frac{1}{2}$, while leaving the 50% penalty rate unchanged, increases both welfare and the present discounted value of lifetime income tax remittances. While there is a trade-off between allowing account holders to keep their money in the tax-benefited account longer and the governmentâs receipt of the tax revenue sooner, our results suggest that the delay in receiving the tax revenue is worth the increase in income tax remittances that come from higher account balances.

Understanding the role the two withdrawal penalties play in savings behavior is increasingly urgent. More than half of U.S. households held at least one IRA, 401(k), or other defined-contribution account by the end of 2017, with $17 trillion of assets saved (Investment Company Institute 2019). Employers have shifted to offering 401(k)s and other defined-contribution accounts instead of defined-benefit accounts such as pensions, and often auto-enroll their employees. States and localities have begun establishing "Auto-IRA" programs in which workers not eligible for other retirement plans are automatically enrolled in an IRA. Contemporaneous to the growing prevalence of these accounts, policy-makers are actively changing the rules around these penalties in response to mounting unease about retirement security.

Previous work studying behavioral responses to these penalties has generally focused on one of the two penalties in isolation, and has not been able to address welfare considerations. In contrast, we consider both penalties simultaneously. This allows us to study the relative magnitudes of their impacts on the timing of withdrawals, and means that our model more accurately captures the trade-offs considered by individuals when they decide how much to contribute or withdraw from these accounts. In addition, our structural model allows us to explicitly consider how changing either penalty would impact welfare, tax remittances, and IRA balances.

We also add to the nascent literature using bunching moments to identify structural parameters via dynamic models. Using bunching moments to identify structural models is a recent methodological development. Our setting differs from previous work using bunching moments to estimate the EIS (Best et al. 2019; Choukhmane 2021). Rather than making a one-time refinancing decision (as in Best et al. (2019)), our individuals make the decision to withdraw, or not, each period. This is similar to Einav et al. (2015), in which individuals choose whether or not to fill a prescription each week based on a health event shock. We estimate an EIS of 1.061. Our EIS estimate is consistent the upper-end of the previous estimates. As IRA holders are higher income individuals than non-IRA holders on average, this result is in line with the argument that, in general, wealthier individuals have a higher EIS (Guvenen 2006).

1 Individual Retirement Accounts

Individual Retirement Accounts (IRAs) are tax-benefited personal savings accounts. IRAs are defined-contribution accounts, meaning that contributions are made into individual accounts. A third of U.S. taxpayers hold at least one IRA.

We focus on traditional IRAs, which comprise about three-quarters of all IRAs. Contributions to traditional IRAs are deducted from taxable income, and withdrawals from traditional IRAs are treated as taxable income at the time of withdrawal.[2] Throughout the paper, we will use the terms withdrawal and distributions to discuss "normal" withdrawals (i.e., we do not include withdrawals due to rollovers or Roth conversions, or the death of the account holder) unless otherwise specified.[3]

Traditional IRAs were introduced as part of the Employee Retirement Income Security Act of 1974 to stimulate retirement saving, particularly among individuals who otherwise would reach retirement with little saved. The key tax benefit is that returns to the principal are not separately taxed. There is no annual tax on accrued interest, and no capital gains tax on withdrawals. Instead, returns are treated as regular taxable income upon withdrawal.[4] There are two additional tax benefits. First, because contributions to a traditional IRA are deducted from taxable income, an account holder's taxable income is lower in a year in which they make a contribution. Second, the amount of income tax ultimately due on contributions to a traditional IRA may be lower if the subsequent withdrawal is taken in a year when the account holder faces a lower marginal tax rate than in the year when she made the contribution. These tax benefits provide a large incentive

[2] In contrast, contributions to Roth IRAs are included in taxable income, but withdrawals from Roth accounts are not subject to federal income taxation. See Appendix A.2 for more details on Roth IRAs.

[3] Funds can be "rolled over" from a different defined-contribution retirement savings account into an IRA. For example, an individual may rollover funds from a 401(k) into an IRA, or from one IRA into another (essentially combining their separate IRA accounts). The vast majority of IRA assets come from rolled over defined-contribution accounts (Goodman et al. 2019). See Appendix A.6 for more details about what happens when an IRA holder dies.

[4] See Appendix A.1 for a simple algebraic example of this benefit.

to contribute to traditional IRAs.

To discourage withdrawing from IRAs before retirement, "early" withdrawals are penalized with an additional 10% tax. This penalty, coupled with a contribution limit ($5,500 in 2015 for individuals under age 50), ensures that the tax benefits of IRAs do not extend to all savings (particularly non-retirement savings).[5] Borrowing from IRAs is not permitted.

To limit the benefits accrued by individuals who would have had sufficient retirement savings without IRAs, a second penalty was introduced: the "excess accumulation penalty." At a certain age, IRA holders are required to take a minimum withdrawal each year. If an IRA holder does not take the required withdrawal, she will owe an additional 50% tax on the difference between the required amount and the withdrawn amount. These required withdrawals, known as Required Minimum Distributions (RMDs), were imposed for IRAs as part of the Tax Reform Act of 1986 (TIGTA 2015-10-042; Mortenson et al. 2019). In addition to ensuring that IRAs do not only amount to a generous tax break (or a tax-benefited bequest), RMDs also minimize the tax revenue cost of IRAs and other defined-contribution retirement savings plans.

We give additional information about the early withdrawal and excess accumulation penalties below. Additional institutional details about traditional IRAs are provided in Appendix A.

1.1 The early withdrawal penalty

Withdrawals from traditional IRAs made before age $59\frac{1}{2}$ face a penalty of 10% on the full amount withdrawn in addition to income tax. For example, if an individual's marginal tax rate was 15% and she made an early withdrawal of $1,000 from a traditional IRA, she would owe both $150 in income tax as well as a $100 penalty (10% of the amount withdrawn early). In 2016, an estimated 1.2 million individual tax returns reported over $1.5 billion in penalties due because of early withdrawals from retirement accounts (Internal Revenue Service 2016). Goodman et al. (2019) estimate that 20% of individual contributions made to IRAs and 401(k)s by individuals below the age of 48 leave the accounts within 8 years of the contribution.

In some circumstances, early withdrawals from an IRA are exempt from the penalty. IRA holders may receive an exemption to pay for qualified medical or higher education expenses. First-time home-buyers may take a withdrawal of up to $10,000 from an IRA without penalty. Individuals who agree to take "substantially equal" withdrawals for a

[5]See Appendix A.4 for more details about IRA contribution limits.

fixed period of time of at least five years may also be exempt from the penalty. Argento et al. (2015) report that, for individuals under age 55, 21% of withdrawals from retirement accounts were penalized.

Previous work on early withdrawals has focused on the circumstances that lead account holders to incur the penalty. This research has consistently found that penalized withdrawals are substantially more likely for individuals that experience adverse shocks such as job loss and divorce, especially among individuals with lower levels of non-retirement assets (see, e.g., Amromin and Smith (2003), Butrica et al. (2010), and Argento et al. (2015)). One recent attempt to quantify the causal effect of these penalties is Goda et al. (2018), who use exact timing of birth to estimate how withdrawals from IRAs respond to the early withdrawal penalty. They find that average withdrawals increase approximately 80% after account holders cross the $59\frac{1}{2}$ threshold, and that the response is largely due to first-time withdrawals.

1.2 The excess accumulation penalty and Required Minimum Distributions

Starting in the tax year that a traditional IRA holder turns a certain age, she is required to take a Required Minimum Distribution (RMD) every year. Because we consider a period before 2019 in our empirical analysis, we use $70\frac{1}{2}$ as the base policy age threshold for RMDs and the excess accumulation penalty throughout the paper.[6] The first RMD tax payment is due by April 1 of the calendar year following the year in which the account holder turns $70\frac{1}{2}$ (i.e., the calendar year in which they turn $71\frac{1}{2}$); subsequent payments must be made by December 31 of each calendar year.[7] More than 15,000 traditional IRA holders reported owing the 50% excess accumulation penalty in 2016, resulting in $7.3 million in additional payments to the IRS (Internal Revenue Service 2016).

The amount of the RMD is based on the balance of the account on December 31 of the previous tax year and life expectancy tables (see Appendix A.5). The penalty for failing to take a required withdrawal is called the "excess accumulation penalty" and is equal to 50% of the required amount not withdrawn. Consider again our individual with a marginal income tax rate of 15%. If she failed to take a required withdrawal of $1,000 from a traditional IRA, she would owe $150 in income tax and a $500 penalty (50% of the money not withdrawn). We discuss how RMDs work in the event of the account holder's

[6] The Setting Every Community Up for Retirement Enhancement (SECURE) Act of 2019 raised the age at which traditional IRA holders are subject to RMDs from $70\frac{1}{2}$ to 72.

[7] If a traditional IRA holder waits to submit her first RMD payment until April 1 of the calendar year after the year in which she turns $70\frac{1}{2}$, she will have two RMDs included in her taxable income in the first year she makes an RMD payment.

death in Appendix A.6.

The empirical literature on RMDs has centered around the "RMD Holiday" in 2009, which suspended the RMD rules for the 2009 tax year. Brown et al. (2017) find that a third of traditional IRA holders subject to RMD rules in 2008 did not take a withdrawal in 2009. Mortenson et al. (2019) use year-over-year variation in addition to the RMD Holiday and estimate that over half of traditional IRA holders are constrained by the RMD rule and would take less than their required withdrawal (or none at all) if not for the RMD rules.

The literature studying the role of these penalties has usually considered them separately. The exception is Sabelhaus (2000), who finds that while these policies change withdrawal timing, total projected taxable withdrawals do not change much if the age for penalty-free withdrawals is lowered to $55\frac{1}{2}$, or the age for required withdrawals is raised to $75\frac{1}{2}$. We take a substantively different approach by estimating a structural model, which allows us to focus on the welfare implications of changing the age thresholds in addition to the behavioral responses. We also consider changes to the penalty rates in addition to changing the age threshold.

2 Administrative tax data

We use de-identified administrative tax data from the Internal Revenue Service (IRS). We create a panel based on a 5% random sample of individuals with Social Security numbers and aged 18 or older in 1999. We follow these individuals through 2015. The 17-year panel is balanced apart from exit due to death and emigration.

We limit our sample to individuals who have an IRA account. We identify individuals as "IRA holders" if we ever observe them making a contribution to an IRA (including rollovers), taking a normal distribution from an IRA,[8] or having an outstanding IRA balance. Our final sample of IRA holders comprises 3,913,401 unique individuals. We focus on traditional IRA holders in our analysis. We observe that 72.2% of IRA accounts are traditional IRAs.[9]

We focus on IRAs because we can cleanly identify contributions to and withdrawals

[8]Normal distributions do not include withdrawals due to rollovers or Roth conversions, or the death of the account holder.

[9]In our sample, 27.8% of IRAs are Roth IRAs. We find that 15.7% of IRA holders hold both kinds of accounts. There are two additional types of IRAs: SIMPLE IRAs and SEP IRAs. These plans make up a small fraction of IRA activity (see Appendix A.3). As these plans face the same tax benefits and penalties as traditional IRAs, we follow Mortenson et al. (2019) and count contributions toward (and withdrawals from) a SEP or SIMPLE IRA as contributions toward (and withdrawals from) a traditional IRA for our analysis.

from IRAs in the administrative tax data, as well as the end-of-year fair market values of IRA accounts. In contrast, we do not separately observe 401(k)s and other employer-sponsored defined-contribution accounts which are subject to the same penalties.[10] Of the $17 trillion dollars saved in IRAs or other defined-contribution accounts in 2017, $9.2 trillion (54%) was in IRAs (Investment Company Institute 2019). There are two facts that suggest our results likely apply to holders of 401(k)s and other employer-sponsored defined-contribution accounts. First, of households that have an IRA, over 80% also have an employer-sponsored retirement plan.[11] Unless households that only have a 401(k) are systematically different from those that have both an IRA and an employer-sponsored account, we will capture behavior from many 401(k) holders in our analysis. Second, many 401(k)s are ultimately rolled over into IRAs when individuals separate from the employers who offered the 401(k).

We supplement our IRA data with information from individual income tax returns such as adjusted gross income, wage and self-employment income, and other sources of retirement income such as Social Security and employer-sponsored defined benefit plans (such as pensions) and defined-contribution plans (such as 401(k)s). We also include data on additional taxes owed due to early withdrawals, non-qualified withdrawals, or failing to take a minimum required withdrawal. We discuss in detail which IRS forms we use, the relevant sample restrictions, and the construction of key variables in Appendix B.

3 Bunching evidence and reduced-form estimates

Measuring bunching around kinks and notches in budget sets has become an increasingly popular tool for estimating behavioral responses to various incentives.[12] This approach was developed by Chetty et al. (2011) for budget set kinks and extended to notches by Kleven and Waseem (2013).

In this section, we adapt the standard approach in order to estimate the number of traditional IRA holders who changed the timing of their withdrawals in response to these penalties, and the amount of money shifted. We estimate four outcomes: the number of individuals who shift the timing of their first withdrawal from a traditional IRA (and the

[10] While information about 401(k)s is captured in survey data, there are limitations to how that data could be used. IRAs and employer-sponsored retirement savings accounts are vastly underreported in both the Survey of Income and Program Participation (SIPP) and the Current Population Survey (CPS) when compared to IRS records (Bee and Mitchell 2017). Larrimore et al. (2019) show that, in 2010, the CPS was missing half a billion dollars in income from pensions, annuities, and IRA withdrawals.

[11] Authors' calculations using Figure 8.7 of Investment Company Institute (2019). Note that, in this figure, employer-sponsored retirement plans include both defined-contribution and defined-benefit plans.

[12] Kleven (2016) provides a recent review of this literature.

amount of money impacted), and the number of individuals who shift the timing of any withdrawal from a traditional IRA (and the amount of money impacted).

3.1 Reduced-form estimation strategy

We fit a flexible polynomial to the empirical distributions of our outcomes. This polynomial excludes data in a region around the age thresholds (the "excluded ranges"), as shown in shown in Equation 1:

$$c_j = \sum_{i=0}^{p} \beta_i (a_j)^i + \sum_{k=a_{59.5,-}}^{a_{59.5,+}} \gamma_k \cdot \mathbb{I}[a_k = a_j] + \sum_{k=a_{70.5,-}}^{a_{70.5,+}} \delta_k \cdot \mathbb{I}[a_k = a_j] + \upsilon_j, \tag{1}$$

where c_j is the value of the outcome at age j, $[a_{59.5,-}, a_{59.5+}]$ is the excluded region around the early withdrawal penalty, $[a_{70.5,-}, a_{70.5+}]$ is the excluded range around the excess accumulation penalty, and p is the order of the polynomial. The β coefficients capture the distribution without the withdrawal penalties (the "counterfactual" distribution), while the γ and δ coefficients represent the difference from the counterfactual distribution brought on by the withdrawal penalties.

We estimate the counterfactual distribution of the outcome using the predicted values from Equation 1, leaving out the portion from the middle terms:[13]

$$\widehat{c_j} = \sum_{i=0}^{p} \widehat{\beta}_i (a_j)^i \tag{2}$$

To determine the excluded range for the early withdrawal penalty, we iterate over all possible combinations of $a_{59.5,-}$ and $a_{59.5+}$ to minimize Equation 3:[14,15]

$$\{a_{59.5,-}, a_{59.5,+}\} = \arg\min \left| \widehat{B}_{59.5} - \widehat{M}_{59.5} \right| \tag{3}$$
$$\widehat{B}_{59.5} = \sum_{j=59.5}^{a_{59.5,+}} (c_j - \widehat{c_j}) \quad \widehat{M}_{59.5} = \sum_{j=a_{59.5,-}}^{58.5} (\widehat{c_j} - c_j)$$

Throughout this section, we will refer to $\widehat{B}_{59.5}$ as the "excess mass" or the "bunching mass,"

[13]We do not allow the counterfactual levels to drop below 0.

[14]This method differs slightly from Kleven and Waseem (2013), who visually determine the equivalent of $a_{59.5+}$ and then iterate over $a_{59.5,-}$ to minimize Equation 3. We chose our method because there was not an obvious visual cue for $a_{59.5+}$. Both approaches assume no extensive margin responses.

[15]All possible values of $a_{59.5,-}$ include $25\frac{1}{2}$ to $58\frac{1}{2}$. All possible values of $a_{59.5,+}$ include $59\frac{1}{2}$ to $69\frac{1}{2}$. Because our bin size is fixed to 1, we are not able to exactly obtain $\widehat{B}_{59.5} = \widehat{M}_{59.5}$. Appendix D shows the value of Equation 3 for all combinations of $a_{59.5,-}$ and $a_{59.5+}$.

and $\widehat{M}_{59.5}$ as the "missing mass." The excess mass is defined as the difference between the empirical and the counterfactual distributions on the low-tax side of the age threshold (in our case, the portion of the excluded range to the right of the age threshold for the early withdrawal penalty). Similarly, the missing mass is the difference between the empirical and the counterfactual distributions on the high-tax side of the age threshold.

The excess accumulation penalty is unusual as a notch because all traditional IRA holders must take withdrawals out at age $70\frac{1}{2}$. This means that, if there are individuals who would rather take their first withdrawal after age $70\frac{1}{2}$, we should see a spike at age $70\frac{1}{2}$ as the age of first withdrawal (which we do), and little to no observations after that (which we also do). This is because we would expect individuals who would rather take the first withdrawal after age $70\frac{1}{2}$ to push the timing of their first withdrawal as late as possible without being penalized. Unlike the case of standard notches, where individuals have an incentive not to cross the threshold, most individuals in our setting will cross the threshold eventually (unless the individual dies before age $70\frac{1}{2}$ or withdraws all of the funds in their IRA before age $70\frac{1}{2}$). The individual has no incentive to deviate from their optimal behavior before age $70\frac{1}{2}$. We can therefore assume that traditional IRA holders who take their first withdrawal in the few years preceding $70\frac{1}{2}$ are behaving optimally. As a result, we simply call age $70\frac{1}{2}$ the excluded range for the excess accumulation penalty.[16] The "excess mass" associated with the excess accumulation penalty, shown in Equation 4, is calculated similarly to that for the early withdrawal penalty, except that we do not estimate an upper bound for the excluded region:[17]

$$\widehat{B}_{70.5} = c_{70.5} - \widehat{c}_{70.5} \tag{4}$$

We estimate the magnitude of the bunching responses using Equation 5, which considers the difference in the cumulative densities above the early withdrawal threshold (below the excess accumulation threshold) compared to the total quantity of the outcome

[16] In other words, $a_{70.5,-} = a_{70.5,+} = 70\frac{1}{2}$.

[17] The spike in the number of individuals who take their first withdrawal at $70\frac{1}{2}$ (and the corresponding total value of those withdrawals) is so large that we are unable to fully distribute that amount under our counterfactual distribution. For the number of individuals withdrawing, the percentage of traditional IRA holders taking their first withdrawal at or above age $70\frac{1}{2}$ that we are not able to account for with the counterfactual distribution is 38.3%. Similarly, the percentage of total first withdrawals taken at or above age $70\frac{1}{2}$ is 3.6%. We believe the tail of the true counterfactual distribution actually extends well past age $80\frac{1}{2}$. We observe age of first withdrawal for Roth IRA accounts well into the 90s. This is consistent with the idea that traditional IRA holders who are forced to take withdrawals from their IRAs because of the RMD rules would otherwise hold on to those funds, allowing them to enjoy more years of tax-benefited growth, until they would then be included in a bequest. See Appendix A.6 for more information about the rules around inheriting an IRA.

in our sample, N:

$$mag_j = \frac{\widehat{B}_j}{N}, \quad j \in \{59.5, 70.5\} \tag{5}$$

mag_j is an estimate for the percentage of the outcome impacted by these thresholds.[18]

We present two pieces of evidence that some traditional IRA holders are responding to these age thresholds. First, we show the distribution of the age at which traditional IRA holders took their first withdrawal, and the total amount withdrawn at each age of first withdrawal. Second, we present the proportion of traditional IRA holders taking withdrawals at each age, and the total amount withdrawn at each age, in a single year. We focus on traditional IRA holders for two reasons: the majority of IRAs are traditional IRAs, and because the tax-benefited in our structural model is designed as a traditional IRA. We present the equivalent figures for Roth accounts in Appendix C.2.

3.2 Bunching by age of first withdrawal

We define the age of first withdrawal as the age at which we first observe an individual making a withdrawal, excluding early withdrawals (before age $59\frac{1}{2}$) with known, qualifying exceptions. This definition is an upper bound on the actual first year of withdrawals, because some individuals may have begun taking withdrawals before our sample period. For this analysis, we exclude individuals who are older than 65 in the first year of our sample so that we do not overestimate the number of individuals taking their "first" withdrawal over age $70\frac{1}{2}$. We also exclude data from 2009. Traditional IRA holders were not subject to the same age thresholds in 2009 as in all other years in our sample as a result of the one-year suspension of the RMD rules included in the Worker, Retiree, and Employer Recovery Act of 2008.

We consider two outcomes by age of first withdrawal: the number of individuals who took their first observed withdrawal at each age, and the amount of money distributed at first withdrawal at each age. Figure 1 shows the empirical distribution (as bars) and our estimated counterfactuals from Equation 1 (as dotted lines) for both outcomes. Figure 1a shows the number of individuals taking their first withdrawal at each age in our sample, while Figure 1b shows the total amount withdrawn by those individuals. In each panel of Figure 1, the two dotted vertical lines show our estimated values for $a_{59.5,-}$ and $a_{59.5,+}$. For these outcomes, we use $p = 6$ to fit the counterfactual.

[18]For example, consider the question of "what percentage of IRA holders changing the timing of their withdrawals?" If we estimated that $\widehat{B}_j = 100$ individuals, and there were $N = 1{,}000$ individuals in the relevant analysis, we would conclude that 10% of individuals were altering their withdrawal timing.

Figure 1: Bunching by age of first withdrawal

(a) Number of IRA holders (b) Gross withdrawal amounts ($bil)

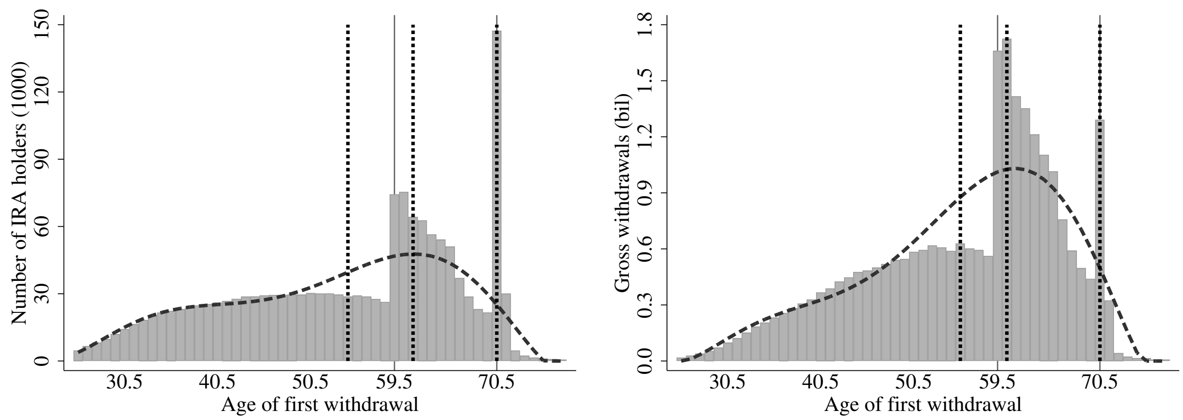

Notes: Figure 1a: $N = 1{,}536{,}392$ unique individuals. Figure 1b based on $25,967,606,263 total withdrawals, inflated to 2015 values. Excludes individuals older than age 65 at the beginning of our sample period, and all data from 2009. We define the age of first withdrawal as the age at which we first observe an individual making a withdrawal, excluding early withdrawals with known, qualifying exceptions. This definition is an upper bound on the actual first year of withdrawals, because some individuals may have begun taking withdrawals before our sample period. Figure 1a shows the number of individuals in our sample we observe taking their first withdrawal from a traditional IRA at each age and the estimated counterfactual distribution. Figure 1b shows the total amount distributed in the first withdrawal by traditional IRA holders who took their first withdrawal at each age and the estimated counterfactual distribution. Withdrawal amounts are inflated to 2015 values. For both figures, the bars show the empirical distribution, and the the dotted lines represent our estimated counterfactual distribution.

It is clear in Figure 1 that traditional IRA holders change their behavior in response to these penalties. There is a dramatic increase in the number of individuals taking their first withdrawal at both age thresholds, and in the amount of money withdrawn in first distributions. The number taking their first withdrawal at age $70\frac{1}{2}$ is significantly higher than the number taking their first withdrawal at age $59\frac{1}{2}$, but the total amount withdrawn at age $59\frac{1}{2}$ is greater than the total amount withdrawn at age $70\frac{1}{2}$. In other words, the average first withdrawal amount for those taking their first withdrawal at age $70\frac{1}{2}$ is considerably less than for those taking their first withdrawal at age $59\frac{1}{2}$, as shown in Figure 2.

There are a small number of individuals whom we observe taking their first withdrawal after age $70\frac{1}{2}$. The vast majority of these cases show the individual taking the withdrawal at age $71\frac{1}{2}$. This is not surprising, given that the first RMD payment is actually due to the IRS by April 1 of the calendar year after the year in which the individual turns $70\frac{1}{2}$ (that is, the calendar year in which they turn $71\frac{1}{2}$).[19] There are two likely explanations for

[19] See Section 1.2 for more details.

Figure 2: Average value of first withdrawals from traditional IRAs ($1,000)

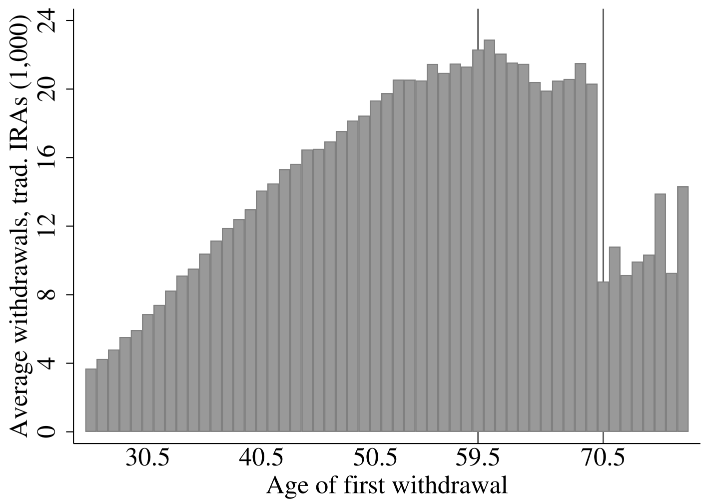

Notes: Based on 1,536,392 unique individuals and $25,967,606,263 total withdrawals, inflated to 2015 values. Excludes individuals older than age 65 at the beginning of our sample period, and all data from 2009. We define the age of first withdrawal as the age at which we first observe an individual making a withdrawal, excluding early withdrawals (before age $59\frac{1}{2}$) with known, qualifying exceptions. This definition is an upper bound on the actual first year of withdrawals, because some individuals may have begun taking withdrawals before our sample period. Figure shows the average amount withdrawn in the first withdrawal from a traditional IRA by c who took their first withdrawal at each age. Withdrawal amounts are inflated to 2015 values.

the small number of traditional IRA holders we observe taking their first withdrawal after age $71\frac{1}{2}$. These individuals may have rolled a 401(k) into an IRA after age $70\frac{1}{2}$. These individuals would not necessarily have held an IRA before the rollover. Another possibility is that these traditional IRA holders may not have fully understood the RMD rules. While these individuals should in theory remit the excess accumulation penalty, the IRS may waive the penalty if the individual can prove that they did not make the required payment as a result of reasonable error and undertook steps to correct their mistake. Mortenson et al. (2019) report that only 2 to 3% of individuals who fail to make an RMD payment violate the RMD rules the following year.

Our estimates for $a_{59.5,-}$, $a_{59.5,+}$, $\widehat{B}_{59.5}$, and $\widehat{B}_{70.5}$, as well as bootstrapped standard errors, are given in Table 1.[20] Table 1 also includes estimates of $mag_{59.5}$, $mag_{70.5}$, and the

[20]Standard errors for $a_{59.5,-}$, $a_{59.5,+}$, $\widehat{B}_{59.5}$, and $\widehat{B}_{70.5}$ are calculated using a bootstrap procedure (see Chetty et al. (2011) and Kleven and Waseem (2013)). We generate alternative data by randomly sampling (with replacement) the residuals produced by Equation 1 and adding those on to the predicted values for each age. Although we force our counterfactual values to be at least 0, we keep the original residual terms from Equation 1. Unlike when calculating the counterfactual distribution, the predicted values in this case are $\widehat{c}_{j,bs} = \sum_{i=0}^{p} \widehat{\beta}_i (a_j)^i + \sum_{i=a_{59.5,-}}^{a_{59.5,+}} \widehat{\gamma}_i \cdot \mathbb{I}[a=i] + \sum_{i=a_{70.5,-}}^{a_{70.5,+}} \widehat{\delta}_i \cdot \mathbb{I}[a=i]$. We then re-calculate $a_{59.5,-}$, $a_{59.5,+}$, $\widehat{B}_{59.5}$, and $\widehat{B}_{70.5}$ using the simulated data as if it were our original data. We repeat this

Table 1: Changes in first withdrawals from traditional IRAs

	Number of taxpayers	Gross withdrawals
Parameter estimates		
N	1,536,392	\$25,967,606,263
$a_{59.5,-}$	54.5	55.5
	(6.2)	(6.1)
$a_{59.5,+}$	61.5	60.5
	(2.8)	(2.8)
$\widehat{B}_{59.5}$	71,962	\$1,345,804,160
	(20,904)	(467,059,441)
$\widehat{B}_{70.5}$	122,088	\$794,694,208
	(5,674)	(139,377,942)
Magnitude estimates		
U.S. total in 2015	201 million	\$277 billion
Proportion holding a traditional IRA	23.6%	n/a
First withdrawal proportion	4.5%	34.7%
Relevant population for scaling	2.1 million	\$94.5 billion
In response to early withdrawal penalty:		
$mag_{59.5}$	4.7%	5.2%
Scaled to U.S., annual	99,566	\$4.9 billion
In response to excess accumulation penalty:		
$mag_{70.5}$	7.9%	3.1%
Scaled to U.S., annual	168,920	\$2.9 billion

Notes: Bootstrapped standard errors given in parentheses. Our preferred estimate of $a_{59.5,+}$ minimize the difference between $\widehat{B}_{59.5}$ and $\widehat{M}_{59.5}$ given $a_{59.5,+} = 49\frac{1}{2}$ (see Appendix D). Scaled amounts are calculated by multiplying the relevant U.S. population in 2015 by the appropriate magnitude estimate. Estimate of number of taxpayers in the U.S. in 2015 based on authors' internal calculations. We multiply this value by the percentage of our sample that hold traditional IRAs to estimate the number of individuals with traditional IRAs, and then by the average proportion of traditional IRA holders that take their first withdrawal each year. Total amount withdrawn from traditional IRAs in 2015 taken from Statistics of Incomes Tax Stats: Accumulation and Distribution of Individual Retirement Arrangements (IRA) Table 1 (2015), available at https://www.irs.gov/pub/irs-soi/15in01ira.xls. We multiply this value by the average amount of withdrawals from traditional IRAs that are part of first withdrawals.

implied number of individuals and amount of money impacted by these age thresholds when scaled to the full population of taxpayers of the United States in 2015.

procedure 1,000 times to obtain a distribution for each estimated variable. We define the standard error for each variable as the standard deviation of the bootstrapped values of that variable. While one benefit of our setting is that we do not have to estimate bin size, the drawback is that we have very few bins (and even fewer after removing the estimated excluded region) and therefore few standard errors to sample for our bootstrap procedure. This means that there are a small number of residuals being used for the bootstrap procedure and, ultimately, to determine our standard errors.

When considering the number of individuals that change when they take their first withdrawal from an IRA, we find $a_{59.5,-} = 54.5$ and $a_{59.5,+} = 61.5$.[21] We estimate that 4.7% of the individuals included in the analysis changed when they took their first withdrawal in response to the early withdrawal penalty, and 7.9% changed when they took their first withdrawal in response to the excess accumulation penalty. These estimates translate into approximately 99,600 and 168,900 individuals changing when they take their first withdrawal due to these penalties each year.

We find similar values for $a_{59.5,-}$ and $a_{59.5,+}$ when we look at the total amount of withdrawn by age of first withdrawal (55.5 and 61.5, respectively). We have more faith in our value for $a_{59.5,-}$ based on the number of individuals rather than the amount of money, because we are not able to identify those exact individuals changing their behavior around first withdrawal (and therefore cannot know precisely which withdrawals were shifted). As this excluded range is smaller than that estimated for the number of individuals who shifted their behavior, the results from this exercise will be conservative relative to the exercise where we used the excluded range estimated for the number of individuals shifting the timing of their first withdrawal.

We estimate that 5.2% of gross first withdrawals are moved as a result of the early withdrawal penalty, and 3.1% of gross first withdrawals are moved as a result of the excess accumulation penalty. If we convert these percentages to annual dollar amounts, we find that $4.9 billion worth of withdrawals is shifted up to seven years as a result of the early withdrawal penalty, and $2.9 billion is withdrawn earlier than would have been without the excess accumulation penalty.

These magnitude estimates should be considered lower bounds on the number of individuals who change their behavior in response to these penalties, and the dollar amount of withdrawals shifted, for four reasons. First, we are focused on traditional IRAs. Roth IRA holders are also subject to the early withdrawal penalty for withdrawals larger than the size of the principal investment. Second, our analysis is focused on individuals who hold IRAs, but the same penalties exist for 401(ks) and other traditional defined-contribution retirement savings accounts. The early withdrawal penalty also applies to a particular type of defined-benefit account: cash balance plans. Third, we focus on the first withdrawal, whereas traditional IRA holders face these thresholds (and corresponding penalties) every age before (after) age $59\frac{1}{2}$ ($70\frac{1}{2}$). The timing of individuals' second, third, fourth, etc. withdrawals could also be impacted by these penalties. Finally, for both populations (IRA and 401(k) holders), these estimates do not take into account the extent to which

[21] The difference between $\widehat{B}_{59.5}$ and $\widehat{M}_{59.5}$ equals 1,412 individuals at these values of $a_{59.5,-}$ and $a_{59.5,-}$.

they affect the extensive margin. In other words, there may be individuals who do not contribute to a retirement savings plan because of these age thresholds.

3.3 Bunching within a single year

In this section, we apply a similar methodology to estimate the impact these penalties have on the withdrawal behavior of traditional IRA holders in the cross-section. This allows us to estimate the magnitude of the effect on all traditional IRA holders, not just those considering their first withdrawal. We focus on a single year to ensure that, as in our previous analysis, each individual appears only once. We picked 2005 because the youngest individuals in our sample were 18 in 1999. We anticipate that the majority who attended college would have graduated within 6 years. We limit this analysis to individuals in our sample who had a non-zero balance in a traditional IRA at the end of 2004 or 2005.

We consider both the proportion of IRA holders who take a withdrawal as well as the average size of withdrawals. Figure 3 shows the empirical distribution and our estimated counterfactuals for both outcomes. The response at the age thresholds is striking. The percentage of account holders who take withdrawals jumps from 10.4% at age $58\frac{1}{2}$ to 20.4% at age $59\frac{1}{2}$ (an increase of 10.0 percentage points). Even more dramatic is the jump at age $70\frac{1}{2}$: the percentage of account holders who take withdrawals increases from from 32.4% at age $69\frac{1}{2}$ to 87.3% at age $70\frac{1}{2}$ (an increase of 54.9 percentage points).[22]

Our cross-sectional analysis differs from that in Section 3.2 in several respects. For these outcomes, we use a fitted first-degree polynomial ($p = 1$) rather than the sixth-order polynomial we used in Section 3.2. This is based on the observation that both trends in Figure 3 appear to be linear before they drop ahead of the early withdrawal penalty. We also exclude the region above $70\frac{1}{2}$ for these estimates. Finally, we visually determine the value of $a_{59.5,-}$ to be $49\frac{1}{2}$ and iterate over possible values of $a_{59.5,+}$ to minimize the difference between $\widehat{B}_{59.5}$ and $\widehat{M}_{59.5}$, rather than iterating over all possible combinations of $a_{59.5,-}$ and $a_{59.5,+}$ as we did in Section 3.2. We make this change because it is visually clear where the proportion of traditional IRA holders taking a withdrawal in the cross section begins to fall ahead of the early withdrawal penalty. The dotted vertical lines in Figure 3 show $a_{59.5,-}$ ($49\frac{1}{2}$ for this analysis) and our estimated values for $a_{59.5,+}$.

[22]There are two likely explanations for why the proportion of account holders taking withdrawals jump to exactly 100% at age $70\frac{1}{2}$. First, our data are pre-audit, and therefore some initial noncompliance could be embedded. Second, it's possible our definition of "who is an IRA holder in 2005" is over-inclusive, and therefore we are including more individuals in the denominator than we should. The withdrawal pattern we observe after age $70\frac{1}{2}$ is very similar to Figure 6 in Mortenson et al. (2019), including the fact that we don't observe 100% of traditional IRA holders taking a withdrawal after age $70\frac{1}{2}$.

Figure 3: IRA withdrawal behavior in 2005

(a) Proportion of trad. IRA holders (b) Average withdrawal amount ($1,000)

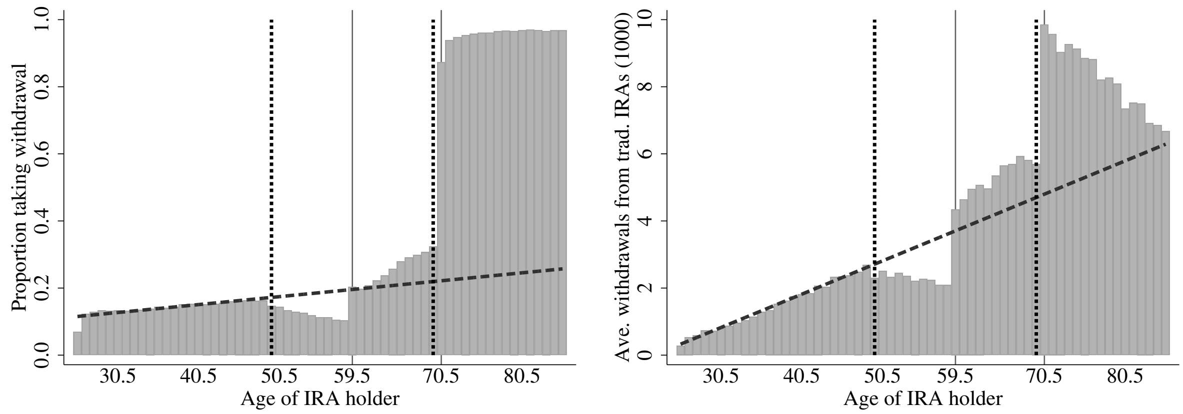

Notes: Figure 3a: $N = 1{,}449{,}868$ unique individuals. Figure 3b: based on $6,483,125,796 total withdrawals, inflated to 2015 values. Excludes early withdrawals (before age $59\frac{1}{2}$) with a known, qualifying exception. Limited to individuals with a positive traditional IRA balance in 2004 or 2005. Figure 3a shows the proportion of account holders taking a withdrawal in 2005 and the estimated counterfactual distribution. Figure 3b shows the average amount withdrawn by age in 2005 and the estimated counterfactual distribution. Average withdrawal amounts are inflated to 2015 values. For both figures, the bars show the empirical distribution, and the the dotted lines represent our estimated counterfactual distribution.

Our estimates for $a_{59.5,+}$ and scaled values for $\widehat{B}_{59.5}$ and $\widehat{B}_{70.5}$, as well as bootstrapped standard errors, are given in Table 2. We find that $a_{59.5,+} = 69\frac{1}{2}$ for both outcomes.[23] Table 2 also includes estimates of $mag_{59.5}$, $mag_{70.5}$, and the implied number of individuals and amount of money impacted by these age thresholds each year when scaled to the full population of taxpayers of the United States in 2015.

We estimate that, in a single year, 1.4% of our sample of traditional IRA holders changes the timing of their withdrawals in response to the early withdrawal penalty. This translates into approximately 648,400 traditional IRA holders in the U.S. changing their withdrawal behavior each year. This is considerably larger than our estimate of the number of individuals who change the timing of their first withdrawal each year (about 99,600), which underscores our point that our estimates based on the timing of first withdrawal are likely lower bounds. We estimate that approximately $12.6 billion is not withdrawn from traditional IRAs each year as a result of the early withdrawal penalty.

Our estimates are even larger for the excess accumulation penalty. We find that 16.2% of our sample of traditional IRA holders change the timing of their withdrawals because of RMDs (about 7.7 million individuals, compared with the 168,900 estimated

[23] Tables 15 and 16 in Appendix D show the difference between $\widehat{B}_{59.5}$ and $\widehat{M}_{59.5}$ for every possible combination of $a_{59.5,-}$ and $a_{59.5,+}$.

Table 2: Changes in withdrawals from traditional IRAs in a single year

	Number of taxpayers	Gross withdrawals
Parameter estimates		
N	1,449,868	$6,483,125,796
$a_{59.5,-}$	49.5	49.5
$a_{59.5,+}$	69.5	69.5
	(0.1)	(0.4)
$\widehat{B}_{59.5}$ (scaled by N)	19,746	$295,103,584
	(2,035)	(29,284,054)
$\widehat{B}_{70.5}$ (scaled by N)	234,677	$1,060,085,440
	(4,080)	(66,206,502)
Magnitude estimates		
U.S. total in 2015	201 million	$277 billion
Proportion holding a traditional IRA	23.6%	n/a
Relevant population for scaling	47.6 million	$277 billion
In response to early withdrawal penalty:		
$mag_{59.5}$	1.4%	4.6%
Scaled to U.S., annual	648,385	$12.6 billion
In response to excess accumulation penalty:		
$mag_{70.5}$	16.2%	16.4%
Scaled to U.S., annual	7,706,011	$45.3 billion

Notes: Bootstrapped standard errors given in parentheses. Our preferred estimate of $a_{59.5,+}$ minimize the difference between $\widehat{B}_{59.5}$ and $\widehat{M}_{59.5}$ given $a_{59.5,+} = 49\frac{1}{2}$ (see Appendix D). Scaled amounts are calculated by multiplying the relevant U.S. population in 2015 by the appropriate magnitude estimate. Estimate of number of taxpayers in the U.S. in 2015 based on authors' internal calculations. We multiply this value by the percentage of our sample that hold traditional IRAs to estimate the number of taxpayers with traditional IRAs. Total amount withdrawn from traditional IRAs in 2015 taken from Statistics of Incomes Tax Stats: Accumulation and Distribution of Individual Retirement Arrangements (IRA) Table 1 (2015), available at https://www.irs.gov/pub/irs-soi/15in01ira.xls.

to change the timing of their first withdrawal each year). We estimate that about $45.3 billion is withdrawn from traditional IRAs earlier than it would have been without the excess accumulation penalty.

The true counterfactual distribution over age $70\frac{1}{2}$ could be higher or lower than what we estimate. Our results for the excess accumulation penalty rely on a strong assumption: that the linear counterfactual distribution is correct past age $70\frac{1}{2}$. If traditional IRA holders become increasingly likely to take withdrawals as they age, our estimates would be lower than the true distribution (i.e., we would overstate the number of individuals impacted). There is evidence that this is true: the apparently linear trend from age $60\frac{1}{2}$ to $69\frac{1}{2}$ has a steeper slope than the linear trend below age $59\frac{1}{2}$. If some individuals who haven't taken a withdrawal by age, say, 75 would never take a withdrawal, the counterfactual distribution would flatten at some point and our counterfactual estimates would be higher than the true distribution (i.e., we would understate the number of individuals impacted).

There are other reasons to believe these estimates are lower bounds. We do not include Roth IRAs, which are also subject to the early withdrawal penalty,[24] or 401(k)s and other traditional defined-contribution retirement savings accounts, which are subject to both of these penalties. As with our estimates based on the age of first observed withdrawal, these estimates do not take into account the extent to which they affect the extensive margin. In other words, there may be individuals who do not contribute to a retirement savings plan as a result of these age thresholds.[25]

3.4 Threats to reduced-form identification

There are two primary threats to identification in our setting. The first threat to identification is the size of the excluded region. The reduced-form bunching approach outlined in this section is less reliable as an estimation strategy when the distortions created by the kink(s) or notch(es) are not very local. While our excluded region is relatively narrow when considering changes to the age of first withdrawal, the excluded range is quite large when we consider the proportion of IRA holders in a single year who are changing the timing of their withdrawals.

The second threat to identification concerns the shape of the estimated counterfactual distribution. A key assumption of our reduced-form analysis is that the counterfactual distribution is smooth through the excluded range. We would worry that the counterfactual distribution might not be smooth if there are other related policies with the same age thresholds. We conduct two diagnostic tests to test the assumption that our observed bunching is specifically due to the age thresholds for penalties related to IRA withdrawals and not other changes in the traditional IRA holders' financial environments. First, we plot the proportion of IRA holders receiving Social Security and receiving a wage, by age. We do not observe bunching at the age threshold for either penalty. Second, we compare the proportion of account holders taking a withdrawal at each age by the half of the year

[24]Figure 15 in Appendix C.2 shows withdrawal behavior in 2005 for Roth IRAs. The proportion of Roth IRA holders taking a withdrawal at any age is always under 0.015, and drops over time. While we observe a small spike at age $59\frac{1}{2}$, the proportion of Roth IRA holders taking a withdrawal continues to fall after that.

[25]It is worth noting that if these penalties resulted in fewer contributions to traditional IRAs, this would (likely) reduce the amount of consumption that is deferred to age $59\frac{1}{2}$ or later.

in which their birthday falls. We observe bunching at ages 59 and 70 for individuals whose half birthday is in the same calendar year as their birthday (e.g., individuals who turn $59\frac{1}{2}$ in the same calendar of the year as 59), but at 60 and 71 for individuals whose half birthday is in the calendar year after their birthday. These diagnostics support our assumption that the bunching we observe is in fact due to the age thresholds related to withdrawal penalties (see Appendix C.1 for more details).

The assumption that the counterfactual distribution is smooth is also problematic if individuals use the age thresholds as reference points. In our setting, this would be true if traditional IRA holders consider the age thresholds as suggestions that they should start taking out withdrawals at those ages. There is a hint of this occurring for Roth accounts. Figure 15a in Appendix C.2 shows a small increase in the proportion of individuals taking withdrawals from Roth IRAs at age $70\frac{1}{2}$ even though the RMD rules do not apply for Roth accounts. In Figure 16, we show that IRA holders with both types of accounts may be more likely to take normal withdrawals from a Roth account at age $70\frac{1}{2}$ than IRA holders who only have a Roth account. That is, we believe individuals who hold both types of accounts may not understand that the RMD rules do not apply to their Roth account as well (or choose to impose a similar heuristic to their Roth account), and therefore begin taking withdrawals at age $70\frac{1}{2}$ from both accounts. The primary concern with reference points in bunching analyses is that it is impossible to disentangle the reference point response from the response to the financial incentive. This means any estimated elasticities will overstate the true structural elasticities. Even if these individuals do use the age thresholds for these penalties as reference points, however, it should not be a problem for our reduced-form results because we estimate the magnitude of the behavioral response rather than underlying behavioral elasticities.

4 The dynamic life-cycle model

Section 3 presented strong evidence of a behavioral reaction to the penalties for early withdrawal and excess accumulation. While our reduced-form results help us answer the question of how many people are shifting the timing of their IRA withdrawals in response to these penalties, we are not able to use our reduced-form estimates to understand the potential welfare and tax revenue impact of changes to these penalties. In order to evaluate counterfactual policies, we develop and estimate a dynamic life-cycle model. The estimated model gives us a framework in which we are able to change these penalties and analyze the impact on savings behavior, welfare, and tax remittances.

The key features of the model are as follows. Every period, individuals choose consumption and how much to save (or dissave) in two different assets: a standard savings account and a tax-benefited account. Dissaving from the tax-benefited account is penalized before period $t < t_e$, and required after period $t > t_{rmd}$. Individuals receive exogenous labor income and, at period t_P, begin receiving an annual pension. Individuals receive utility from consumption and from a bequest motive.

Section 4.1 provides details about the model, and Section 4.2 sets up the individual's problem. A complete list of the parameters used in the model is given in Appendix F.

4.1 Model set-up

Lifespan and survival probabilities Individuals live for no more than T periods. The conditional probability of living to period t conditional on surviving to period $t-1$ is π_t. We adopt the norm that the t subscript refers to the beginning of the period.

Exogenous income The log of labor income y_t is determined by two components: a deterministic component that is a function of the period, and a stochastic component:

$$\ln y_t = g(t, X) + \varepsilon_t^y. \tag{6}$$

The stochastic component, ε_t^y, follows an AR(1) process with normally distributed errors with mean 0 and variance σ_ε^2:

$$\begin{aligned} \varepsilon_t^y &= \eta \varepsilon_{t-1}^y + \zeta_t \\ \zeta &\sim \mathcal{N}\big(0, \sigma_\zeta^2\big) \end{aligned} \tag{7}$$

Formally modelling labor supply decisions is outside the scope of the model. Instead, individuals receive exogenous labor income in each period, and decreasing labor supply over the lifespan is captured in decreased labor income. Individuals receive a Social Security-style pension at an exogenously given age, t_p, which is known in advance. The value of the pension income is a function of the period at which the individual begins to receive it. Total income, z_t, is the sum of labor income and the annual pension:

$$z_t = y_t + P \cdot \mathbb{I}[t \geq t_P] \tag{8}$$

Savings Individuals have access to two types of savings accounts: a standard savings account (S) and a tax-benefited account (A). At the beginning of period t, the stock of assets in the standard savings account is denoted as S_t and in the tax-benefited account as A_t. The individual chooses how much to (dis)save in both accounts: s_t in the standard savings account and a_t in the tax-benefited account. $s_t > 0$ indicates savings; $s_t < 0$ indicates withdrawals. The same is true for a_t.

The two accounts follow similar laws of motion, which are summarized in Equations 9 and 10. After the individual has chosen her (dis)saving amounts s_t and a_t for period t, the amount is added from the account balance at the beginning of the period. The stock of assets in each type of savings account at the start of the next period is equal to the balance at the end of the previous period after interest. The two savings accounts face different pre-tax rates of return, with $r_S < r_A$. There is no borrowing.

$$S_{t+1} = (1 + r_S)(S_t + s_t) \tag{9}$$

$$A_{t+1} = (1 + r_A)(A_t + a_t) \tag{10}$$

Contributions to the tax-benefited account are deducted from taxable income, while withdrawals from the tax-benefited account are added to taxable income. Contributions are capped at a period-dependent level, $\overline{a}_t$. Withdrawals from the tax-benefited account before period t_e are penalized at rate p_e. Minimum withdrawals, $a_{t,rmd}$, are required after period t_{rmd}, with failure to withdraw penalized at rate p_{rmd}. The cost of saving in the tax-benefited account is reduced liquidity before period t_e. This is relevant given uncertainty in the earnings process.

The individual faces two potential penalties for withdrawals from the tax-benefited account. The early withdrawal penalty is given by τ_e in Equation 11. If $a_t < 0$ (indicating a withdrawal) and $t < t_e$, the individual owes an additional $p_e \cdot |a_t|$ in taxes.

$$\tau_e = p_e \cdot |a_t| \cdot \mathbb{I}[a_t < 0, t < t_e] \tag{11}$$

Starting in period t_{rmd}, individuals are required to take Required Minimum Distributions from their tax-benefited savings account. Required Minimum Distributions are a function of the beginning-of-period account balance and the period:

$$a_{t,rmd} = f(A_t, t)$$

Failure to take the minimum amount required triggers the excess accumulation penalty, represented by τ_{rmd} in Equation 12. If $|a_t| < a_{t,rmd}$ and $t \geq t_{rmd}$, the individual owes an additional $p_{rmd} \cdot (a_{t,rmd} - |a_t|)$ in taxes.

$$\tau_{rmd} = p_{rmd} \cdot (a_{t,rmd} - |a_t|) \cdot \mathbb{I}[a_t < a_{t,rmd}, t \geq t_{rmd}] \tag{12}$$

Taxation The individual owes income tax in each period. Income tax owed is determined by a function $\tau_y(\cdot)$ with its argument equal to the sum of total income z_t and withdrawals from the tax-benefited savings account. Total taxes owed are equal to income tax plus any penalties due to early withdrawals or failure to take a Required Minimum Distribution:

$$T(z_t, a_t, \tau_e, \tau_{rmd}) = \tau_y\big(z_t - a_t\big) + \tau_e + \tau_{rmd}. \tag{13}$$

Utility from consumption Utility from consumption is given by a constant EIS function:

$$u(c_t) = \frac{\sigma}{\sigma - 1} c_t^{\frac{\sigma-1}{\sigma}}, \tag{14}$$

where σ is the elasticity of intertemporal substitution.

Bequest motive Individuals value end-of-life wealth W via the warm-glow bequest motive given in Equation 15:

$$B(W_{t+1}) = \theta \frac{\alpha}{\alpha - 1} \left(\frac{W_{t+1}}{\theta} \right)^{\frac{\alpha-1}{\alpha}}, \tag{15}$$

where $W_{t+1} = (1 + r_S)S_{t+1} + (1 + r_A)A_{t+1}$. That is, if an individual dies between periods t and $t + 1$, the value of her bequest is equal to what would have been her total starting wealth in period $t+1$. The specification of the bequest motive follows that used in Jakobsen et al. (2019). The individual does not owe any taxes on $B(W_{t+1})$.

4.2 The individual's maximization problem

An individual starts period t knowing the following state variables: the period (t), her exogenous labor income shock (ε_t^y), the level of assets in both savings accounts (S_t and A_t), the period when she will receive her annual pension (t_P), and what will be the value of her annual pension (P). We collectively describe these state variables as $\Omega_t = \{t,\ \varepsilon_t^y,\ S_t,\ A_t,\ t_P,\ P\}$.

Knowing the state variables, individuals make two choices each period to maximize expected lifetime utility: how much to (dis)save in both savings accounts (s_t and a_t). This decision implies post-tax consumption (c_t) and what would be left in a bequest.

The individual's problem in recursive form is to pick $\{c_t, s_t, a_t\}$ to maximize the Bellman equation given by Equation 16:

$$V_t(\Omega_t) = u(c_t) + \beta\Big(\pi_{t+1}\mathbb{E}\big[V_{t+1}(\Omega_{t+1})\big] + (1-\pi_{t+1})B(W_{t+1})\Big) \tag{16}$$

subject to the following constraints:

$$\text{Budget constraint: } z_t + S_t + A_t = c_t + T(z_t, a_t, \tau_e, \tau_{rmd}) + \frac{S_{t+1}}{1+r_S} + \frac{A_{t+1}}{1+r_A} \tag{17}$$

$$\text{No borrowing constraint: } S_{t+1} \geq 0, \quad A_{t+1} \geq 0 \tag{18}$$

as well as the savings laws of motion (Equations 9 and 10).[26]

V_t is the present value of expected lifetime utility at period t. The model is governed by four preference parameters: β is the discount factor, σ is the elasticity of intertemporal substitution, α is the bequest elasticity, and θ is the weight put on the bequest motive. Individuals save both to protect against bad labor income draws and to finance additional consumption later in life when labor income is low. The exogenous discontinuities in the budget constraint caused by the two penalties mean there is not a closed form solution to the model. We provide details on how we solved the problem numerically in Appendix G.

5 Model results

We estimate the model in two steps. In the first step, we set the value of parameters that can be cleanly estimated outside of the model (e.g., survival probabilities) or are institutional in nature (e.g., the income tax schedule). In the second step, we use the Simulated Method of Moments (SMM) to estimate the preference parameters of the model: the elasticity of intertemporal substitution (EIS), the discount factor, and the two parameters governing the bequest motive. This two-step process is standard in the estimation of life-cycle models (see, e.g., French (2005)).

We present the first-step parameter estimates in Section 5.1 and the second-step parameter estimates in Section 5.2. We describe the estimation procedure in Section 5.2.1, the variation that identifies the four preference parameters in Section 5.2.2, the preference

[26] In general, individuals are not allowed to borrow against IRAs. In 2020, the Coronavirus Aid, Relief, and Economic Security (CARES) Act allowed individuals to take early withdrawals without penalty to pay for specific COVID-19-related expenses, and to repay the withdrawal if the repayment is within three years of the withdrawal.

parameter estimates in 5.2.3, and model fit in Section 5.2.4.

5.1 First step estimates

5.1.1 Survival probabilities and tax policy parameters

Lifespan and survival probabilities We model individuals starting at age $t = 40\frac{1}{2}$ and set $T = 85\frac{1}{2}$.[27] Survival probabilities for each age are calculated using the U.S. Social Security Actuarial Life Tables for 2010. We normalize the survival parameters so that the probability of death at age $85\frac{1}{2}$ is 1.[28] The probabilities are given separately for men and women; we use the average of the two.

Tax-benefited savings account We set the parameters specifying the age and penalty levels for the early withdrawal and excess accumulation penalties equal to the statutory levels for traditional IRAs: withdrawals before age $59\frac{1}{2}$ are subject to an additional $p_e = 10\%$ tax due on the amount withdrawn, while minimum withdrawals must be taken after age $70\frac{1}{2}$ or be subject to an additional penalty of $p_{rmd} = 50\%$ on the amount not withdrawn. The required distribution schedule is modeled after the true schedule and discussed in Appendix A.5.

Contributions to the tax-benefited savings account are capped. Because all dollar values in the model are inflated to their 2015 equivalents, we use the contribution limits from 2015: the contribution limit was $5,500 for individuals under age 50 and $6,500 for individuals aged 50 or older. For more on contribution limits, see Appendix A.

Income tax schedule Income is taxed according to the U.S. 2015 single-filer income tax brackets, given in Table 3. We assign all individuals the standard deduction value for single-filers in 2015 ($6,300).

Table 3: U.S. 2015 single-filer income tax schedule

For income over...	Marginal tax rate	For income over...	Marginal tax rate
$0	10.0%	$189,300	33.0%
$9,225	15.0%	$411,500	35.0%
$37,450	25.0%	$413,200	39.6%
$90,750	28.0%		

[27] In the 2010 U.S. Social Security Actuarial Life Tables, about 10% of individuals survive past age 85.

[28] See Table 6 of Life Tables for the United States Social Security Area 1900-2100, available at https://www.ssa.gov/OACT/NOTES/as120/LifeTables_Tbl_6_2010.html. We use the survival probability to age 85 for age $85\frac{1}{2}$, the survival probability to age 84 for age $84\frac{1}{2}$, etc.

5.1.2 First step parameters estimated from the data

We describe below how we determine initial asset holdings, effective rates of return to the two savings accounts, the labor income process parameters, and pension receipt in the structural model. We make three sample restrictions to our main panel to generate the sample used to estimate the first step parameters (hereafter referred to as the "model input sample"). These restrictions ensure that the model input sample and our model set-up are internally consistent. First, we only include individuals who were never married during our sample period. This means that all of the first step parameters are estimated on individuals rather than households. Second, we limit the sample to individuals who had non-missing income in all periods. Third, we keep individuals who were born between 1920 and 1968. This restricts the sample to individuals who were between the ages of 30 and 79 in 1999. This yields a balanced panel with 60,615 unique individuals.

Initial asset holdings To determine initial balances in both the regular savings account and the tax-benefited savings account, we fit lognormal distributions to the empirical distributions at age 40 using maximum likelihood estimation. We estimate the initial empirical distribution of the regular savings account balance using information from IRS Form 1099-INT. We use the fair market value of end-of-year traditional IRA balances reported on Form 5498, winsorized at the 1st and 99th percentiles, to obtain the empirical distribution of the tax-benefited savings account. We find that the mean of the lognormal distribution of initial tax-benefited assets is lower than that of regular savings (5.433 versus 9.277), but that the standard deviation is larger (3.530 versus 1.765). Our simulated individuals receive one draw from the distribution for regular savings and one from the distribution for tax-benefited savings.[29]

Rates of return We calibrate the values of $1 + r_S$ and $1 + r_A$ to be 1.02 and 1.05, respectively. Having a sufficient wedge in the rate of return for the standard savings account and tax-benefited saving account is critical for the model simulations to fit the data. Investments in tax-benefited accounts enjoy abnormally high real after-tax rates of return for multiple reasons. The predominant reason is that returns to tax-benefited accounts are not subject to annual taxation. Instead, returns accumulate tax-free and are subject to tax upon withdrawal, a treatment that effectively permits investors to obtain compounded returns on deferred tax liabilities. Any positive inflation rate increases nominal returns and affects the effective tax rate on investments outside of tax-preferred accounts, further increasing the difference between real after-tax rates of return in taxable

[29]See Appendix B.1 for additional details about the initial distribution of the regular savings account.

and tax-benefited accounts (Feldstein 1976; Feldstein et al. 1978). In addition, IRAs tend to be heavily invested in equities, whereas assets in non-IRAs are often more diversified. The equity premium is a long-standing puzzle in economics (see, e.g., Mehra and Prescott (1985); Benartzi and Thaler (1995)).

We estimate $1 + r_A$ directly from our data. We limit the model input sample to account-years when the account holder did not not take a withdrawal or make a contribution or rollover, and when the observed growth was between 0% and 15%. After accounting for inflation, the average observed growth rate is 5.1%. The value of $1 + r_S = 1.02$ corresponds to the long-term real interest rate assumed by the 2020 OASDI Trustees Report.[30]

Labor income We estimate the deterministic component of exogenous labor income $g(t, X)$ using a fixed effects model that is cubic in age with year-of-birth fixed effects, as shown in Equation 19. Labor income is defined as the sum of wage income and self-employment income.

$$\ln y_{it} = \rho_0 + \rho_1 t + \rho_2 t^2 + \rho_3 t^3 + YOB_i + \varepsilon_{it}^y \tag{19}$$

The constant is estimated as the average of the year-of-birth fixed effects. We set the fixed effect equal to 0 in our simulations as there is only one cohort.

The residuals from Equation 19 follow an AR(1) process, as given in Equation 20:

$$\begin{aligned} \varepsilon_{it}^y &= \eta \varepsilon_{it-1}^y + \zeta_{it} \\ \zeta_{it} &\sim \mathcal{N}\left(0, \sigma_\zeta^2\right) \end{aligned} \tag{20}$$

We estimate the labor income process using the sum of wage income and self-employment income. The results of this exercise are given in Table 4:

Table 4: Exogenous labor income process parameter estimates

Deterministic component				AR(1) Parameters	
$\hat{\rho}_0$	$\hat{\rho}_1$	$\hat{\rho}_2 * 100$	$\hat{\rho}_3 * 1000$	$\hat{\eta}$	$\hat{\sigma}_\zeta^2$
7.682	0.143	-0.185	0.0036	0.829	0.267

Notes: N = 60,615.

Our estimate of the persistence parameter η is within the range of previous estimates.

[30] The long-range annual real interest rate assumed by the 2020 Old Age, Survivors, and Disability Insurance (OASDI) report was 1.8% in the high-cost scenario, 2.3% in the intermediate-cost scenario, and 2.8% in the low-cost scenario. As the assumed real interest rates used in the annual OASDI Trustees report have been slowly falling over the past decade, we err on the low side and use 2% (OASDI 2020).

For example, Laibson et al. (1998) finds a persistence parameter of 0.511 for high school drop-outs and 0.688 and 0.686 for high-school graduates and college graduates, respectively, while Laibson et al. (2018) estimates a persistence parameter of 0.782 and Choukhmane (2021) reports a persistence parameter of 0.974.

Annual pension Individuals are exogenously assigned the age at which they begin receiving the Social Security-style pension. The probability of claiming the pension at a given age is determined by the proportion of individuals in our data that we observe first receiving Social Security at each age between 62 and 70. Simulated individuals receive the average value of (non-zero) Social Security received by the individuals in the model input sample whom we first observe claiming Social Security at that age.[31] Figure 4 summarizes the probability that an individual claims the pension at each age (shown by the solid lines) and the value of the pension when claimed at each age (given by the bars).

Figure 4: Value of annual pension and probability of claiming

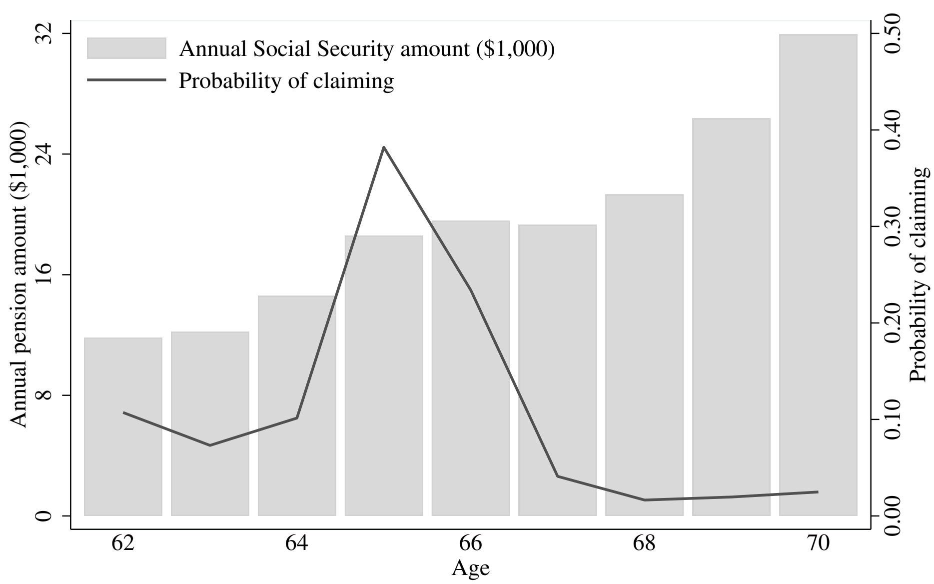

Notes: N = 3,460 individuals. Includes individuals aged 55-60 in 1999 whom we observe first receiving non-zero Social Security between the ages of 62 and 70, inclusive. The probability that an individual claimed at each age is given by the solid lines. The average amount of Social Security given at each age is shown by the bars. Social Security amounts are inflated to 2015 values.

[31] We calculate these averages using the amount received one year older than the claiming age to ensure that we capture a full year of Social Security receipt.

5.2 Second step preference parameter estimates

5.2.1 Estimation procedure: Simulated Method of Moments

As our dynamic model does not have a closed-form solution, we solve the model computationally using the Simulated Method of Moments (SMM). The estimation routine is as follows. We set an initial guess for the four preference parameters, Φ_0: the elasticity of intertemporal substitution (σ_0), the discount factor (β_0), the bequest elasticity (α_0), and the weight on the bequest motive (σ_0). Using backwards iteration, we solve the model using Φ_0 and the output from the first step estimates described in Section 5.1.

We use the model solution to generate S sets of simulated data. For each simulated dataset $m_s(\Phi)$ ($s = \{1, 2, ..., S\}$), we calculate a series of moments denoted $h(m_s(\Phi))$. Note that the simulated data, and therefore the value of $h(m_s(\Phi))$, are a function of the structural parameters Φ used to solve the model.

We compare the average of these simulated moment vectors to the values of the same moments calculated from the administrative tax data, $h(w)$ (where w represents the true empirical data). Specifically, we calculate Equation 21:

$$g(w, \Phi) = h(w) - \frac{1}{S} \sum_{s=1}^{S} h(m_s(\Phi)). \tag{21}$$

Our SMM estimator $\hat{\Phi}_{SMM}$ is then defined as the solution to the following minimization problem:

$$\hat{\Phi}_{SMM} = \arg\min_{\Phi} Z(\Phi, n) \equiv g(w, \Phi)' \hat{W}_n g(w, \Phi), \tag{22}$$

where $\hat{W}_n$ is a positive definite weighting matrix. We set $\hat{W}_n$ to be the identity matrix. Using the identity matrix places more weight on the moments that are largest in absolute value, whereas the optimal weighting matrix (i.e., the inverse of the variance-covariance matrix of the empirically estimated moments) places more weight on the moments that are most precisely estimated. As is shown in Table 6, some of our moments are estimated with a considerably larger sample size than others, which mechanically makes them more precisely estimated. This results in a model fit that quantitatively distorts our counterfactual exercises by providing a better fit for some age groups than for others. As our moments are generally of the same magnitude, we find that using the identity matrix yields better model fit overall relative to the optimal weighting matrix.

We minimize $Z(\Phi, n)$ using Nelder-Mead optimization. We generate 10,000 individuals for each simulation. More details about the numerical solution method are given in Appendix G.

5.2.2 Estimation moments and identification

We jointly estimate the four preference parameters of the model, all of which govern savings decisions. Rather than rely solely on the savings profile to separately identify these parameters, we take advantage of the bunching moments generated by the early withdrawal and excess accumulation penalties. Matching on bunching and other quasi-experimental moments provides more credible variation for identifying preference parameters in structural models.

We use 25 moments in our estimation procedure: the average value of IRA withdrawals at each age between $50\frac{1}{2}$ and $74\frac{1}{2}$. We show in a simple, highly-stylized, model that bunching in our setting can be used to identify the EIS and the discount factor in Appendix E. In our model, individuals face uncertain survival and receive utility from the ability to leave a bequest. The expected utility from leaving a bequest impacts consumption and, subsequently, withdrawal decisions. As a result, average withdrawal sizes enable us to identify the two bequest parameters of the model.

5.2.3 Preference parameter estimates

The preference parameter estimates are given in Table 5:

Table 5: Preference parameter estimates

σ: EIS	β: Discount factor	A: Bequest weight	α: Bequest elasticity	χ^2
1.061	0.920	2.660	1.069	307
(0.157)	(0.00857)	(0.461)	(0.157)	d.f. = 307

Notes: This table show the estimated preference parameters of the structural model described in Section 4. Standard errors are shown in parentheses. The optimal weight matrix used to calculate the standard errors is generated via bootstrap.

We estimate an EIS of 1.061. This is considerably higher than that estimated by both Best et al. (2019) and Choukhmane (2021), who estimate the EIS equal to about 0.1 and 0.4, respectively, but not inconsistent with the upper-end of the range of estimates found in the literature.[32] Chetty (2006) notes that, in models where the EIS is equal to the

[32] For example, Gourinchas and Parker (2002) report a value for the coefficient of relative risk aversion (CRRA) of 0.514 when using their robust weighting matrix, which implies an EIS of $1/0.514 = 1.95$ in

Figure 5: Average adjusted gross income for IRA and non-IRA holders, by age

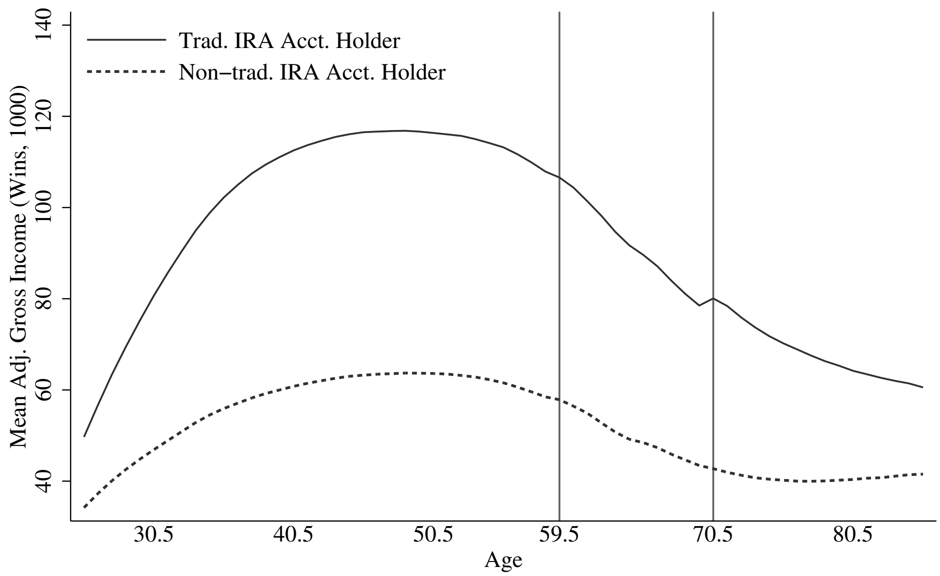

Notes: N = 60,615 unique individuals. Adjusted gross income winsorized at the 1st and 99th percentiles. Adjusted gross income amounts are inflated to 2015 values.

inverse of the CRRA, it is unreasonable to find an EIS less than 1, as it would require an uncompensated wage elasticity lower than any value previously estimated. Guvenen (2006) argues that, in general, wealthier individuals have a higher EIS. Figure 5 compares average adjusted gross income for IRA holders and non-IRA holders at each age. IRA holders have considerably higher average income at every age, implying that they are very likely wealthier and one would expect a higher EIS estimate for this population.

The model is formally rejected by a χ^2 overidentification test. This is primarily driven by the fact that the estimation procedure does not incorporate the variance of the parameters estimated in the first step (e.g., the income process parameters).[33]

We present t-statistics on the moment conditions in Table 6. We find large t-statistics on many of the moment conditions. This is in part the result of the fact that the model is highly stylized, and that many of the moments are estimated on samples with thousands (if not tens of thousands) of observations.

their framework. Best et al. (2019) use Epstein-Zin-Weil preferences to separately estimate the EIS and the CRRA, which could explain some of the disparity in our estimates.

[33]Our standard errors are likely understated for the same reason.

Table 6: Empirical moments, simulated moments, and t-statistics

Age	Empirical moment	Simulated moment	t-stat	Age	Empirical moment	Simulated moment	t-stat
50	1,578	860	13.97	63	4,654	4,502	1.75
51	1,620	937	13.00	64	4,525	5,325	−4.11
52	1,731	1,031	13.37	65	4,595	6,293	−8.90
53	1,721	1,161	11.18	66	4,306	2,393	10.27
54	1,747	1,316	8.93	67	3,847	2,873	6.92
55	1,876	1,508	9.09	68	3,542	3,461	0.96
56	2,022	1,725	6.41	69	3,613	3,992	−0.95
57	1,972	1,976	3.81	70	6,994	7,436	−1.34
58	1,794	2,287	−0.56	71	8,025	7,489	2.13
59	3,356	3,297	1.03	72	7,583	7,562	0.26
60	4,265	3,919	5.10	73	7,787	7,500	0.40
61	4,438	4,673	−0.40	74	7,670	7,454	0.32
62	4,697	3,888	5.72				

Notes: Empirical moments show the average withdrawal value at that age, conditional on having a non-0 account balance in the IRA, as estimated in the model input sample. The simulated moments are the average withdrawal value at that age, conditional on having a non-0 account balance in the IRA, as estimated in our simulated data (N = 10,000).

5.2.4 Model fit for matched moments

Figure 6 compares the simulated data moments to the empirical moments. The figure shows the average withdrawal from a traditional IRA taken at each age, conditional on having a non-0 traditional IRA balance at the beginning of the period. The solid lines represent the moments from the empirical data; the dashed lines show the moments from the simulated data. The shaded regions include moments that are not targeted in the estimation procedure. The key feature of Figure 6 is that we accurately capture the bunching at ages $59\frac{1}{2}$ and $70\frac{1}{2}$. This is important because our counterfactual policy analysis will involve changing the two tax penalties that cause the observed bunching.

The dips in simulated average withdrawals between ages of $60\frac{1}{2}$ and $70\frac{1}{2}$ correspond with ages when the majority of our simulated individuals begin claiming their annual pension. Just over 10% of our sample starts claiming their pension at age $62\frac{1}{2}$, after which we observe a dip in average IRA withdrawals at that age (see Figure 4 in Section 5.1). By age $66\frac{1}{2}$, 90% of our sample are receiving their annual pension, which aligns with the second dip in average withdrawal value.

Figure 6: Comparing matched moments from the empirical and simulated data

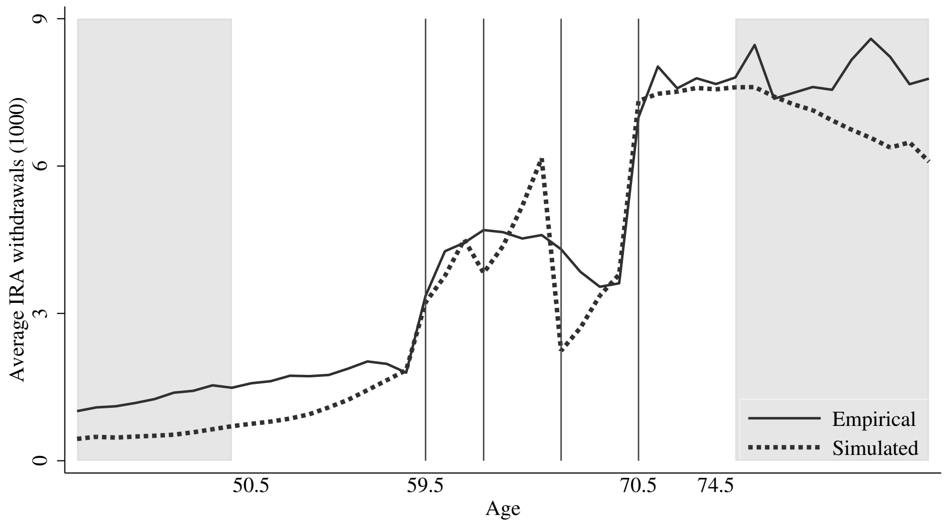

Notes: $N = 60{,}615$ unique individuals. We simulated 10,000 individuals. The figure shows the average withdrawal from a traditional IRA taken at each age, conditional on having a non-0 IRA balance at the beginning of the period. The solid lines represent the moments from the empirical data; the dashed lines show the moments from the simulated data. The shaded regions include moments that are not targeted in the estimation procedure.

5.2.5 Model fit for unmatched moments

We compare the empirical and simulated moments for two additional measures of IRA behavior to further validate our model: average balance by age, and the proportion of account holders taking a withdrawal at each age. We consider the averages for each value conditional on having a non-0 traditional IRA balance at the beginning of the period.

Figure 7 shows the results of this exercise. The solid lines represent the moments from the empirical data; the dashed lines show the moments from the simulated data. Figure 7a compares average balance values. We match the empirical moments well until age $59\frac{1}{2}$, after which we mirror the shape but are a bit lower in levels. This is in part due to the fact that we do not have rollovers in the model, which would pull up average empirical account balances. After age $70\frac{1}{2}$, average balances are influenced by the value of the required minimum withdrawals, which are themselves a function of average balances. It's not a surprise that our average simulated values remain below the average empirical values in this region.

Figure 7b shows the proportion of individuals taking a withdrawal at each age. We are able to capture the magnitude of bunching even in this out of sample exercise.

Figure 7: Comparing unmatched moments from the empirical and simulated data

(a) Average IRA balances ($1,000), by age (b) Proportion taking withdrawal, by age

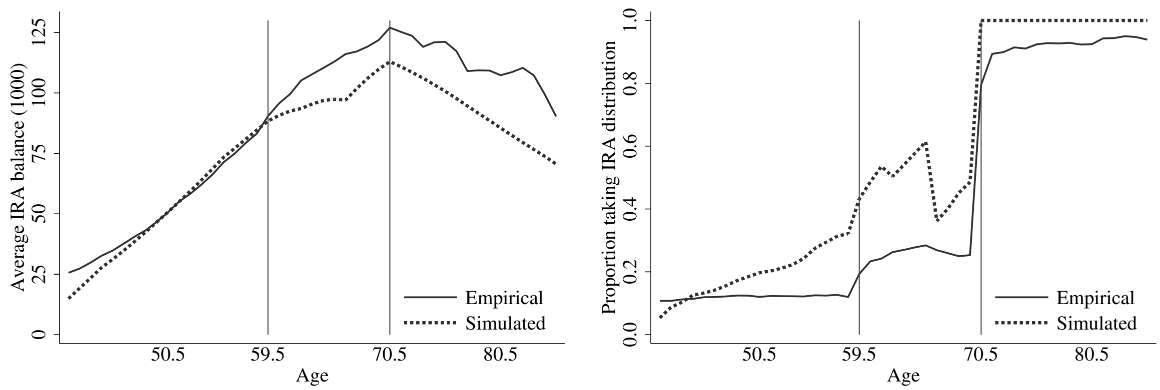

Notes: $N = 60{,}615$ unique individuals. We simulate 10,000 individuals. Panel (a) shows the average IRA balance, by age. Panel (b) shows the proportion of traditional IRA holders taking a withdrawal, by age. The solid lines represent the moments from the empirical data; the dashed lines show the moments from the simulated data.

6 Counterfactual policy analysis

These penalties have been changed to achieve policy goals. For example, the SECURE Act increased the age at which RMDs kick in from $70\frac{1}{2}$ to 72. This change was made with little understanding about what the long term impact would be on welfare and tax remittances. In fact, there is no evidence about whether these penalties do or do not achieve their purported goals of increasing retirement consumption while limiting the tax cost. In this section, we estimate the impact of the change to the age threshold for RMDs from the SECURE Act before considering more broadly a large number of combinations of age thresholds and penalty rates.[34]

6.1 Increasing the age for Required Minimum Distributions from $70\frac{1}{2}$ to $72\frac{1}{2}$

The SECURE Act raised the RMD age to 72, allowing for longer accumulation in tax-benefited accounts. We examine the impact on IRA withdrawal behavior after changing the age for RMDs from $70\frac{1}{2}$ to $72\frac{1}{2}$.[35] Figure 8 shows the changes in average withdrawal

[34]Lawmakers have also changed these penalties during the last two major financial crises. RMDs were suspended in 2009 as part of the Worker, Retiree, and Employer Recovery Act of 2008. In March 2020, the Coronavirus Aid, Relief, and Economic Security Act (CARES Act) waived the early withdrawal penalty for withdrawals due to COVID-19-related financial hardships and suspended RMDs for 2020.

[35]We solve the model in whole periods and therefore can only consider changes in whole years. Because we consider a slightly bigger change than what was actually enacted in 2019, our estimates should be considered upper bounds.

Figure 8: Impact of raising the age for Required Minimum Distributions from $70\frac{1}{2}$ to $72\frac{1}{2}$ on average IRA withdrawals and consumption

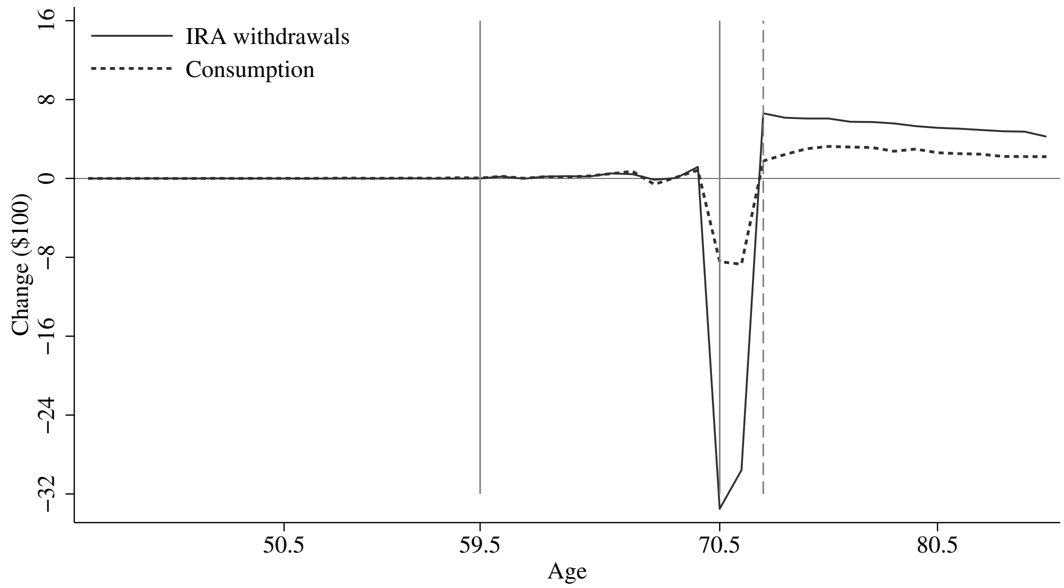

Notes: Each iteration includes 10,000 unique simulated individuals. The figure shows the difference in average IRA withdrawals and consumption from the base policy after we raise the age for Required Minimum Distributions from $70\frac{1}{2}$ to $72\frac{1}{2}$. A positive number means the value is greater under the counterfactual policy; a negative number means the value is lower under the counterfactual policy. The solid lines show average IRA withdrawals, and the dashed lines show average consumption levels.

amounts and consumption at each age relative to the current policy. A positive number means the value is greater under the counterfactual policy; a negative number means the value is lower under the counterfactual policy. The solid lines show average IRA withdrawals, and the dashed lines show average consumption levels.

We observe that average IRA withdrawals decrease by just over \$3,000 at age $70\frac{1}{2}$ after the age for RMDs increases. The decrease at age $71\frac{1}{2}$ is nearly as large. Average consumption also drops, though by less than \$1,000 at both $70\frac{1}{2}$ and $71\frac{1}{2}$. When individuals are required to take withdrawals from their account, they can choose to use the withdrawal to finance consumption or save the withdrawal at the standard savings rate. In contrast, when individuals are not required to take withdrawals, their decision is to take a withdrawal for consumption or continue to save that money in the tax-preferred account. Our results imply that the optimal consumption choice under those two scenarios is not the same.

Average IRA withdrawals are both higher at age $72\frac{1}{2}$ relative to the current policy, and at every age thereafter. This must be true mechanically if average IRA balances are higher due to decreased withdrawals at ages $71\frac{1}{2}$ and $72\frac{1}{2}$. We also see elevated levels of consumption at age $72\frac{1}{2}$, and for the rest of the simulated lifespan, relative to the level under the current policy.

There is essentially no change in withdrawals before age $65\frac{1}{2}$. There is similarly no change in consumption, suggesting that increasing this tax benefit may not induce people to increase their contributions to IRAs from the base level. This means that the additional tax benefit may not be sufficient to encourage individuals to save more in IRAs at the expense of current consumption or other types of precautionary savings. We observe slight changes in both average withdrawals and consumption between ages $65\frac{1}{2}$ and $70\frac{1}{2}$.

We compare four additional outcomes under the current policy and after raising the age for Required Minimum Distributions to $72\frac{1}{2}$. The first is a measure of equivalent variation, which we use to estimate the welfare consequences of raising the age for required withdrawals. To estimate this measure, we calculate the average present discounted value of lifetime utility under both the base policy and the counterfactual policy as shown in Equation 23. Let u_{it} be the utility in period t for individual i under the original policy and u'_{it} be the utility in period t under the counterfactual policy. The present discounted value of lifetime utility for individual i under the base policy is calculated by discounting the value of u_{it} back to age $41\frac{1}{2}$, and then summing up between the ages for $41\frac{1}{2}$ and $85\frac{1}{2}$.[36] We then take the average of this value to determine the average present discounted value of lifetime utility. The average present discounted value of lifetime utility under the counterfactual policy is calculated similarly.

$$u_{PDV,base} = \frac{1}{N}\sum_{i=1}^{N}\sum_{t=41.5}^{85.5}\beta^{t-41.5}(u_{it}) \tag{23}$$

$$u_{PDV,cf} = \frac{1}{N}\sum_{i=1}^{N}\sum_{t=41.5}^{85.5}\beta^{t-41.5}(u'_{it})$$

Once we have determined $u_{PDV,base}$ and $u_{PDV,cf}$, we use our policy functions from solving the model under the base policy to determine how much income we would need to give to (or take away from) our simulated individuals in the first period so that $u_{PDV,base} = u_{PDV,cf}$. A positive number indicates that our simulated individuals are better off under the counterfactual policy, while a negative number indicates that our simulated individuals are better off under the base policy.

In addition to our measure of equivalent variation, we also calculate the difference in the present discounted value of lifetime tax remittances,[37] average IRA balance at age

[36] We discount to age $41\frac{1}{2}$ rather than age $40\frac{1}{2}$ because starting at age $41\frac{1}{2}$, less than 2% of our simulated individuals are constrained by our highest grid value. See Appendix G for more details.

[37] We use Equation 23 to calculate the present discounted value of total taxes remitted, with one change:

$65\frac{1}{2}$, and average bequeathed IRA balance (i.e., average IRA balance at age $85\frac{1}{2}$).

Table 7 shows the results of these four outcomes. As with our measure of equivalent variation, a positive number indicates a higher value in the counterfactual policy, while a negative number indicates a lower value in the counterfactual policy. We find that increasing the age for Required Minimum Distributions from $70\frac{1}{2}$ to $72\frac{1}{2}$ yields marginally higher welfare and slightly lower tax remittances. Bequeathed IRA balances are more than 6% larger after the change. The average IRA balance at age $65\frac{1}{2}$ is slightly lower, which is consistent with what we observed in Figure 8 (i.e., that average withdrawals at age $65\frac{1}{2}$ are higher under the counterfactual policy).

Table 7: Comparing the base policy to the SECURE Act

	Equivalent variation	PDV lifetime tax remittances	IRA balance at age $65\frac{1}{2}$	Bequeathed IRA balance
Levels ($)	339	-81	-186	4,283
Percent change	0.0017%	-0.0254%	-0.1908%	6.6%

Notes: Each iteration includes 10,000 unique simulated individuals. The table shows (1) our measure of equivalent variation, (2) the present discounted value of lifetime taxes remitted, (3) the value of average IRA balances at age $65\frac{1}{2}$, and (4) the value of the bequeathed IRA at age $85\frac{1}{2}$ for the base policy and the counterfactual policy where we raise the age for Required Minimum Distributions from $70\frac{1}{2}$ to $72\frac{1}{2}$.

As shown in Table 8, the decrease in tax revenue is driven by reduced excess accumulation penalty payments. As we'll discuss in Section 6.2, there are good reasons to believe that the relevant measure with the excess accumulation penalty is actually the change in income taxes remitted. In that case, we observe that both welfare and tax remittances increase relative to the base policy in the world where we raise the age for required withdrawals to $72\frac{1}{2}$.

Although we do not account for redistribution in our model, the fact that this policy increases both welfare and lifetime income tax remittances suggests that there are changes to these tax penalties that could yield better outcomes than what we have now. At the very least, we could give individuals a lump sum transfer in the last period of their lives (thereby avoiding potential distortionary effects from redistribution), and they would be strictly better off. In the next two sections, we calculate our measure of equivalent variation and the impact on lifetime tax remittances and IRA balances for a large number of changes to both tax penalties. For these estimates, we change both the penalty rate, and the age at which the penalty applies.

instead of discounting by the discount factor β, we discount by the rate of return for the non-tax-benefited savings account: $\frac{1}{1.01}$, where 1% is an estimate of the government's borrowing rate.

Table 8: Changes in sources of tax revenue from raising the age for Required Minimum Distributions from $70\frac{1}{2}$ to $72\frac{1}{2}$

Average total tax revenue remitted between ages $41\frac{1}{2}$-$85\frac{1}{2}$	
Original policy	317,800.75
RMD age = $72\frac{1}{2}$	317,720.09
Total change	-80.66
Change in income taxes	27.25
Change in early withdrawal penalties	-0.26
Change in excess accumulation penalties	-107.61

Notes: Each iteration includes 10,000 unique simulated individuals. This table shows the average total taxes remitted under the original policy and under the counterfactual policy, the difference between the total, and how changes in the three sources of tax revenue contribute to that difference.

6.2 Changing the excess accumulation penalty

In the previous section, we considered one specific change to the early accumulation penalty: changing the age at which the excess accumulation penalty applies from $70\frac{1}{2}$ to $72\frac{1}{2}$. We considered this particular change because it most closely mirrored the change implemented in 2019, but there is no reason to believe that this change led us to the optimal policy, conditional on the existence of these penalties.

To consider a more exhaustive set of counterfactual policies, we compare the base policy for the excess accumulation penalty (age threshold = $70\frac{1}{2}$, penalty rate = 50%) with every combination of age thresholds $68\frac{1}{2}$ and $78\frac{1}{2}$ and penalty rates 40% to 70% (in increments of 10%). The results are given in Figure 9; Table 18 in Appendix H gives these results both in levels (as shown below) and as percent changes from the base policy.

Figure 9a gives the results for equivalent variation. Each point shows the amount of additional income we would need to give our simulated individuals for the average present discounted value of lifetime utility in the base policy to be equal to that in the counterfactual policy. The straight horizontal line at 0 indicates where there is no difference from the base policy; the base policy is indicated by the black square on the 0 line. Points above the 0 line indicate a positive value relative to the base policy (i.e., the individual is better off under the counterfactual policy), while points below the 0 line indicate a negative value relative to the base policy (i.e., the individual is worse off under the counterfactual policy).

Raising the age for required withdrawals increases welfare. This is not surprising, because raising the age is equivalent to loosening a constraint and we are modeling rational agents with exponential discounting. We find a positive value for all penalty rates considered for every age threshold above the base policy. The observed increases are small in magnitude, with the largest change being an increase of income of just under \$600. The rate of change of the equivalent variation value decreases as the age threshold rises. All of the observed differences result from changing the age threshold; there is essentially no difference if we change the penalty rate.

While there is no explicit government budget constraint in the model, we can consider whether or not the policy changes we consider would be tax-revenue neutral (or tax-revenue increasing). The results for the present discounted value of total lifetime tax remittances are shown in Figure 9b. Raising the age for required withdrawals has a mixed impact on total lifetime tax remittances. Across penalty rates, total lifetime tax remittances increase from age $68\frac{1}{2}$ to $69\frac{1}{2}$, but then decrease as the age threshold rises. Total tax remittances are higher relative to the base policy at ages $68\frac{1}{2}$ through $72\frac{1}{2}$ for a penalty rate of 70%, at ages $68\frac{1}{2}$ through $71\frac{1}{2}$ for a penalty rate of 60%, and at ages $68\frac{1}{2}$ and $69\frac{1}{2}$ for a penalty rate of 50%. As with equivalent variation, the changes are modest in magnitude, with the largest increase falling just over \$100.

The decreases in total tax remittances are largely driven by changes in remittances of the excess accumulation penalty. Mortenson et al. (2019) find that less than 1% of traditional IRA holders who do not comply with the RMD rules file the required Form 5329 to remit the excess accumulation penalty. As most individuals do comply with the RMD rules, this does not amount to a large number of individuals not paying the penalty. It does, however, suggest that including changes in excess accumulation penalty payments may not be the most accurate measure of changes in tax remittances if we change the age threshold and penalty rate for the excess accumulation penalty, because our simulated individuals are forced to pay the penalty.[38]

Figure 9c shows the change in total lifetime income tax remittances. In this case, we observe positive differences from the base policy at all age thresholds through for $76\frac{1}{2}$ a penalty rate of 70%, and at all age thresholds except for $75\frac{1}{2}$ for a penalty rate of 60%. At a penalty rate of 50%, there is an increase in income tax remittances at age $69\frac{1}{2}$ and at ages $71\frac{1}{2}$ through $73\frac{1}{2}$.

All of the counterfactual policies with age thresholds at $71\frac{1}{2}$ or higher that resulted in higher total lifetime tax remittances also had positive equivalent variation values. This

[38]Over 75% of our simulated individual-years have no excess accumulation penalty. All of our simulated individual-years with a positive excess accumulation penalty value took a positive withdrawal. About 75% of those face the penalty on a value that is less than 10% of the amount withdrawn. This suggests our overestimate of the amount of excess accumulation penalties is in part due to simulation error.

means that there are numerous combinations of age threshold and penalty rate that are both welfare improving while also resulting in increased tax revenue.

An important caveat to the results for taxes remitted is that we do not account for taxes paid on inherited IRAs. The rules for what happens to an inherited IRA are complicated. The amount of tax due, and when, depends on who is the beneficiary (e.g., spouse or non-spouse), the age of the beneficiary, and whether or not the original account holder had started taking RMDs (see Appendix A.6 for more details). While modeling these rules is not the goal of this project, we can roughly estimate whether taxes collected on inherited IRAs will increase or decrease simply by looking at whether average IRA account balances increase or decrease in the final period relative to the base policy.

Figure 9d shows the difference in the amount chosen to leave in the IRA as part of the bequest in the last period relative to the base policy. Inherited IRAs increase in size as the penalty rate decreases for a given age threshold, and increase in size as the age threshold increases for a given penalty rate. As with equivalent variation and tax remittances, the variation in age threshold has a much bigger impact on the average inherited IRA balance than variation in the penalty rate. Most importantly, average bequeathed IRA balances are higher for nearly every combination of age threshold and penalty rate that also yielded increased welfare and income tax remittances (the exception being an age threshold of $71\frac{1}{2}$ with a penalty rate of 70%).

6.3 Changing the early withdrawal penalty

We conduct the same exercise for the early withdrawal penalty. We compare our outcomes under the base policy to every combination of age threshold between $55\frac{1}{2}$ and $65\frac{1}{2}$ and penalty rates of 5%, 10%, 20%, and 30%. The results are given in Figure 10; Table 19 in Appendix H gives these results both in levels (as shown below) and as percent changes from the base policy. As with Figure 9, the straight horizontal line at 0 indicates where there is no difference from the base policy; the base policy is marked by the black square on the 0 line. Points above the 0 line indicate a positive amount relative to the base policy, while points below the 0 line indicate a negative amount relative to the base policy.

As shown in Figure 10a, our equivalent variation measure decreases both as the age threshold increases and as the penalty rate increases. This follows from the fact that, for rational agents with exponential discounting, raising either the age threshold or the penalty rate tightens the lifetime budget constraint. As a result, our simulated individuals ought to be worse off in expectation. Unlike the excess accumulation penalty, where varying the age threshold had the biggest impact on welfare, changing the age threshold for the early

Figure 9: Changing the excess accumulation penalty

(a) Change in age $40\frac{1}{2}$ income in base policy to get PDV lifetime utility in counterfactual policy

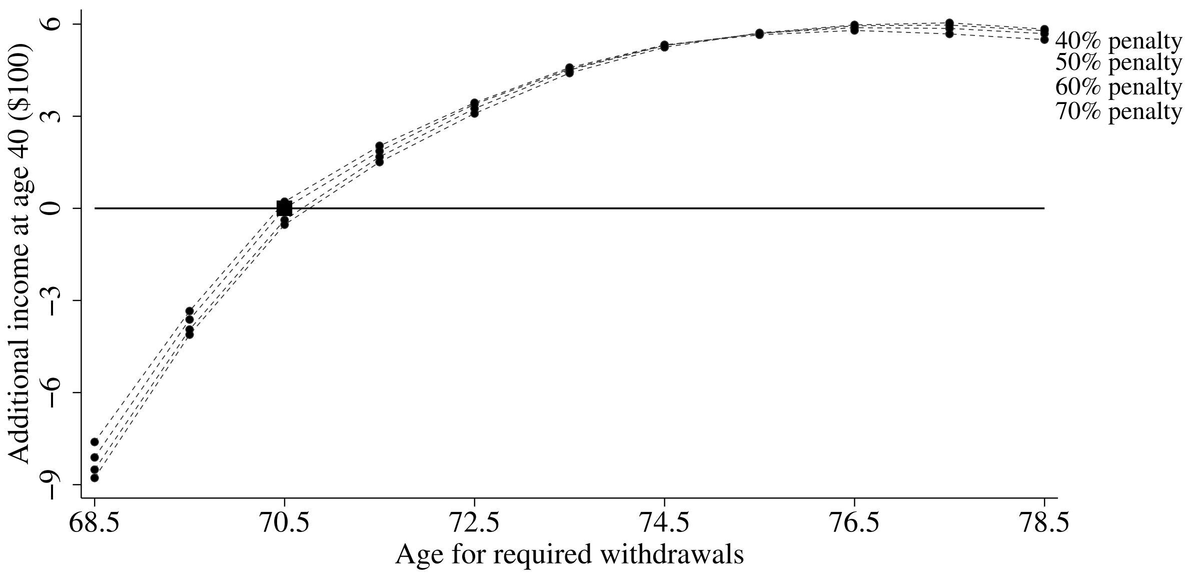

(b) Difference in lifetime taxes remitted relative to base policy ($, PDV from age $41\frac{1}{2}$)

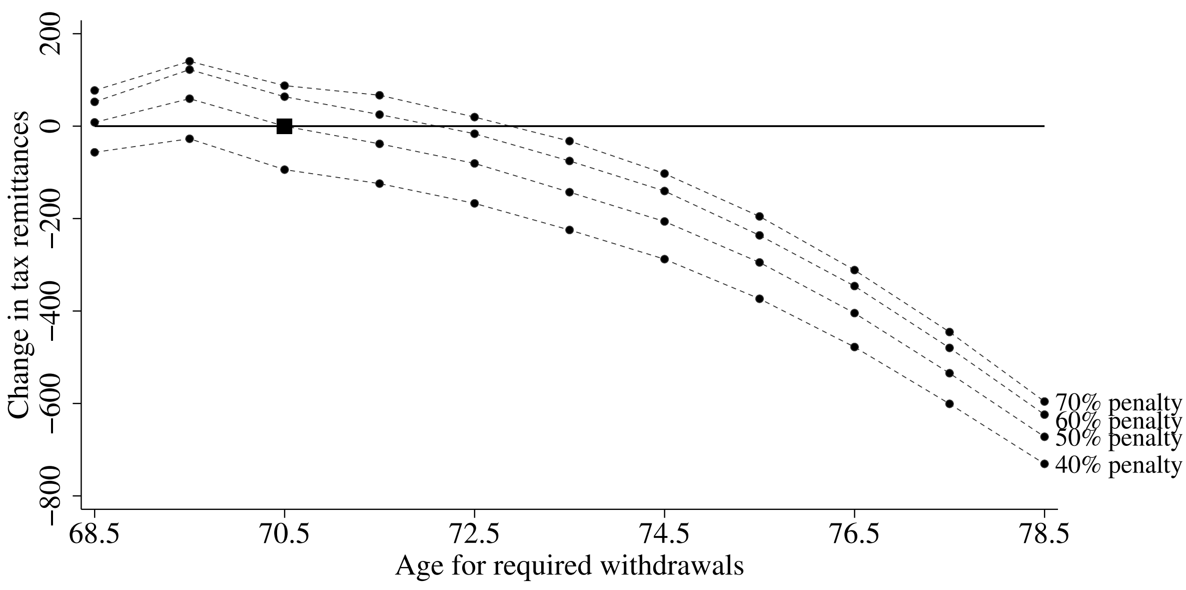

Notes: Each iteration includes 10,000 unique simulated individuals. The figure compares various counterfactual policies for the excess accumulation penalty to the base policy (age threshold = $70\frac{1}{2}$, penalty rate = 50%). Figure 9a shows our measure of equivalent variation. Figure 9b shows the change in the PDV of lifetime tax remittances. Figure 9c shows the change in the PDV of lifetime income tax remittances. Figure 9d shows the change in bequeathed IRA balances. In all panels, the base policy is indicated by a black square on the 0 line.

Figure 9: Changing the excess accumulation penalty, continued

(c) Difference in lifetime income taxes remitted relative to base policy ($, PDV from age $41\frac{1}{2}$)

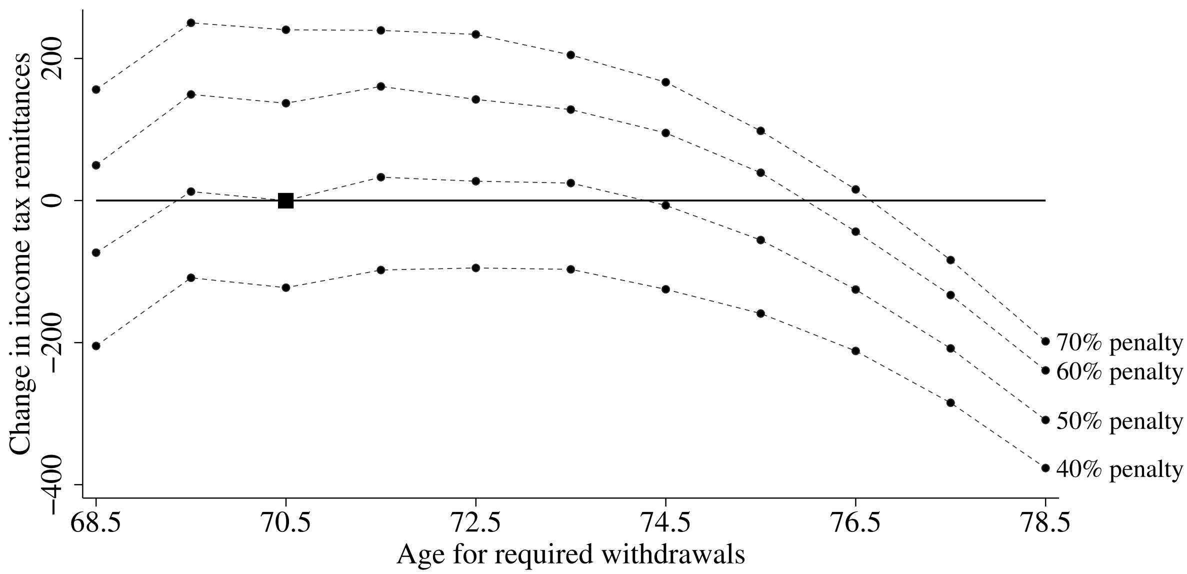

(d) Difference in bequeathed IRA balance relative to base policy

Notes: Each iteration includes 10,000 unique simulated individuals. The figure compares various counterfactual policies for the excess accumulation penalty to the base policy (age threshold = $70\frac{1}{2}$, penalty rate = 50%). Figure 9a shows our measure of equivalent variation. Figure 9b shows the change in the PDV of lifetime tax remittances. Figure 9c shows the change in the PDV of lifetime income tax remittances. Figure 9d shows the change in bequeathed IRA balances. In all panels, the base policy is indicated by a black square on the 0 line.

withdrawal penalty has a smaller impact on welfare than changing the penalty rate.

The impact of changing the early withdrawal penalty on lifetime tax remittances is shown in Figure 10b. Increasing the penalty rate and the age threshold both lead to increased lifetime tax remittances. Increasing the age threshold has an increasingly large impact on lifetime tax remittances.

There are four age threshold-penalty rate combinations that lead to increases in both welfare and lifetime tax remittances: a penalty rate of 5% with an age threshold between $62\frac{1}{2}$ and $65\frac{1}{2}$. The increases in both measures are modest: the equivalent variation measure ranges from \$1,486 to \$2,004, while the change in lifetime tax remittances ranges from \$41 to \$412.

We present the results for lifetime income tax remittances in Figure 10c. We observe that income tax remittances decrease with age threshold, and with the penalty rate. There are increased income tax remittances relative to the base policy at all age thresholds facing a 5% penalty, and for all age thresholds below $59\frac{1}{2}$ facing a 10% penalty. These are the same age threshold and penalty rate combinations that led to increased welfare. As we observe individuals paying the early withdrawal penalty in the tax data, we believe that the correct measure of changes in tax remittances are total changes in tax remittances when considering the early withdrawal penalty.

Because the purported goal of the early withdrawal penalty is to improve financial security later in life, we consider the change in IRA balances at age $65\frac{1}{2}$ in Figure 10d. IRA balances at age $65\frac{1}{2}$ decrease with age thresholds until an age threshold of $58\frac{1}{2}$, and then increase until an age threshold of $65\frac{1}{2}$. Generally speaking, a higher penalty rate results in a higher IRA balance at age $65\frac{1}{2}$, though the higher the age threshold, the bigger the difference between penalty rates. There are numerous combinations of age threshold and penalty rate that lead to higher IRA balances at age $65\frac{1}{2}$ relative to the base policy. Most relevant to us, all three of our counterfactual policies that lead to increased welfare and total tax remittances also result in higher IRA balances at age $65\frac{1}{2}$.

6.4 Discussion of counterfactual exercises

We study how welfare and tax remittances change if we alter the age threshold and penalty rate for the excess accumulation penalty and the early withdrawal penalty. We find that increasing the age for Required Minimum Distributions from $70\frac{1}{2}$ to be between $71\frac{1}{2}$ and $73\frac{1}{2}$, with a penalty rate of 50% or higher, increases both welfare and the present discounted value of lifetime income tax remittances, suggesting that the increase in income taxes remitted from higher account balances may be worth the delayed tax revenue. There may be additional increases in taxes that we do not account for as a result of taxes remitted

Figure 10: Changing the early withdrawal penalty

(a) Change in age $40\frac{1}{2}$ income in base policy to get PDV lifetime utility in counterfactual policy

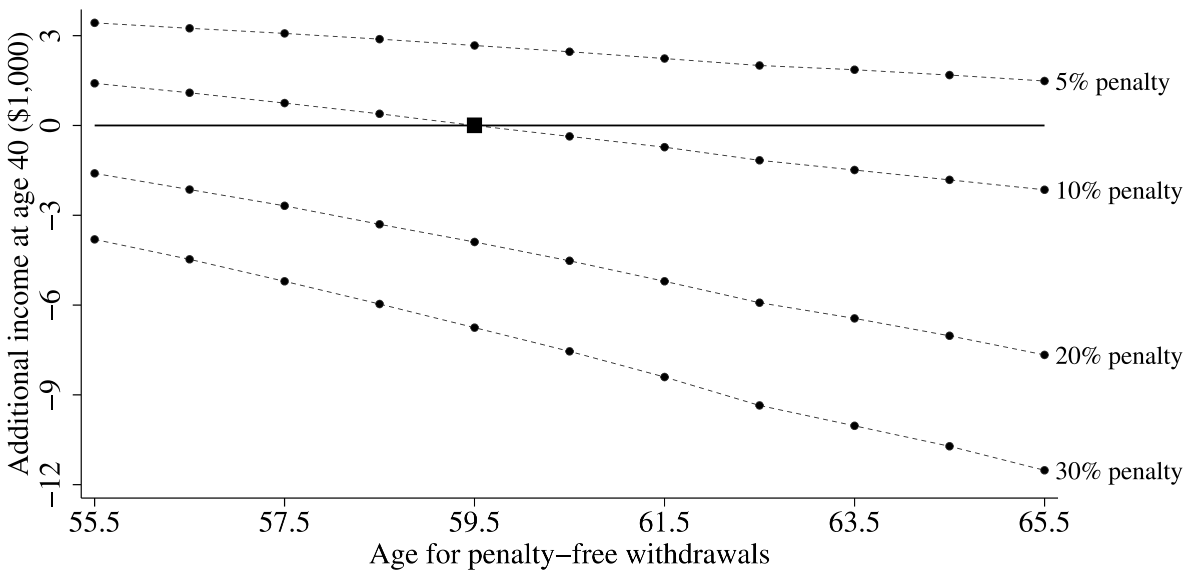

(b) Difference in lifetime taxes remitted relative to base policy ($, PDV from age $41\frac{1}{2}$)

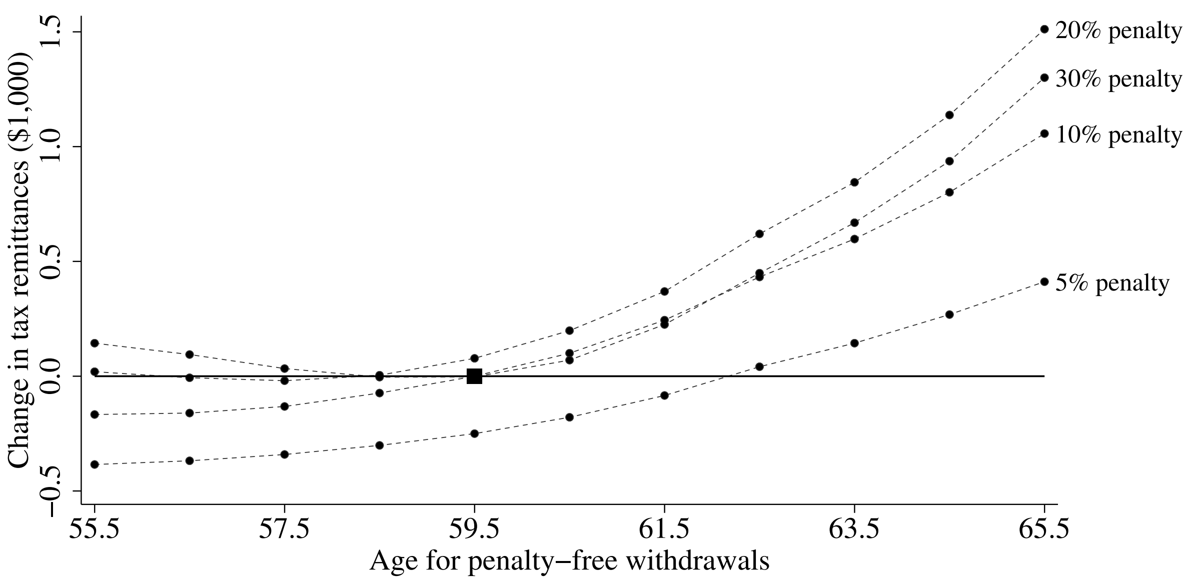

Notes: Each iteration includes 10,000 unique simulated individuals. The figure compares various counterfactual policies for the early withdrawal penalty to the base policy (age threshold = $59\frac{1}{2}$, penalty rate = 10%). Figure 10a shows our measure of equivalent variation. Figure 10b shows the change in the PDV of lifetime tax remittances. Figure 10c shows the change in the PDV of lifetime income tax remittances. Figure 10d shows the change in IRA balances at age $65\frac{1}{2}$. In all panels, the base policy is indicated by a black square on the 0 line.

Figure 10: Changing the early withdrawal penalty, continued

(c) Difference in lifetime income taxes remitted relative to base policy ($, PDV from age $41\frac{1}{2}$)

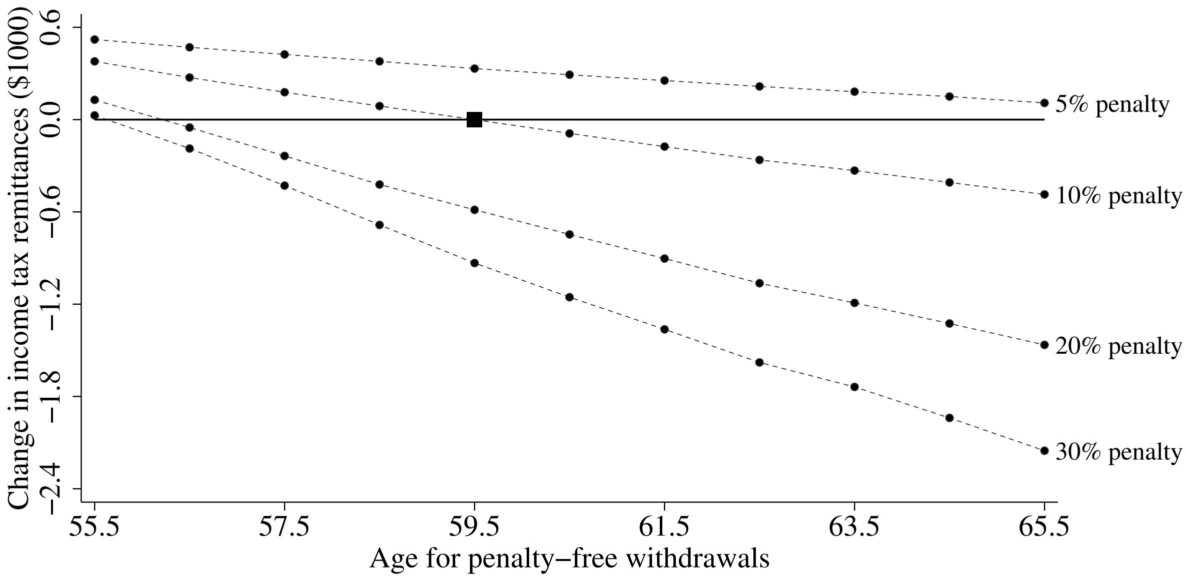

(d) Difference in IRA balances at age $65\frac{1}{2}$ relative to base policy

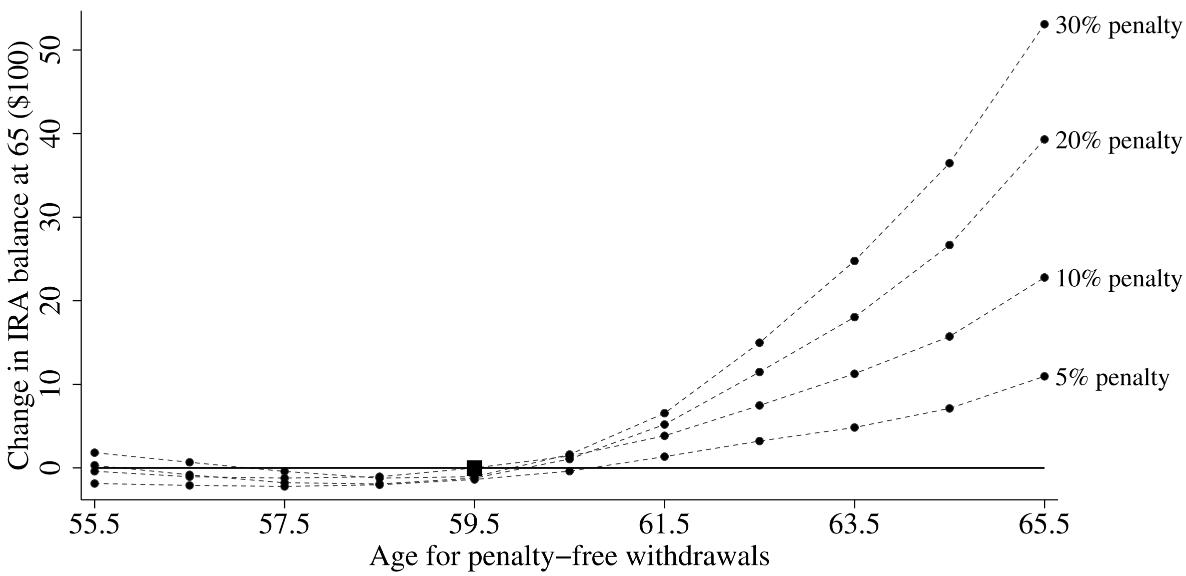

Notes: Each iteration includes 10,000 unique simulated individuals. The figure compares various counterfactual policies for the early withdrawal penalty to the base policy (age threshold = $59\frac{1}{2}$, penalty rate = 10%). Figure 10a shows our measure of equivalent variation. Figure 10b shows the change in the PDV of lifetime tax remittances. Figure 10c shows the change in the PDV of lifetime income tax remittances. Figure 10d shows the change in IRA balances at age $65\frac{1}{2}$. In all panels, the base policy is indicated by a black square on the 0 line.

when IRAs are inherited, because bequeathed IRA balances also increase when the age for Required Minimum Distributions increases.

We find that increasing the age for penalty-free withdrawals to be between $62\frac{1}{2}$ and $65\frac{1}{2}$, while lowering the penalty rate to 5%, increases both welfare and lifetime total tax revenue. Average IRA balances at age $65\frac{1}{2}$ are also higher at these age thresholds and penalty rates. The intuition for this finding is as follows: conditional on having these penalties in place, there are benefits to encouraging individuals to keep their money in these accounts as long as possible. However, if individuals need to take early withdrawals as the result of an unexpected income shock, they can do so with minimal sanction.

There are four things worth mentioning when considering these results. First, we have only changed one policy at a time. Our analysis of changes to the excess accumulation penalty were done keeping the base policy for the early withdrawal as is, and vice versa. There may be policies that involve changes to both penalties that yield even bigger changes in welfare and lifetime tax remittances.

Second, we have only altered two of the policy levers associated with IRAs: the two withdrawal penalties. There is a third policy lever that we have not changed: the contribution limit. The contribution limit increased substantially over the course of our sample period, from $2,000 in 1999 to $5,500 in 2015. Increasing the contribution limit would allow individuals to grow their IRA balances faster, but at the cost of putting more of their savings out of reach in the event of an unexpected income shock. Considering the role of the contribution limit is a direction for future work.

Third, these are partial equilibrium effects. Our model does not account for redistribution of tax revenue. Redistributing tax revenue would have an income effect, which could impact consumption and savings decisions. We also do not model endogenous labor supply choices. Reaching age $59\frac{1}{2}$ does not change our simulated individuals' labor income, but it's possible that individuals reduce their labor supply in response to having penalty-free access to savings in their IRAs. There could also be general equilibrium shifts in the interest rate as a result of policy-induced changes in the supply of savings.

Finally, we do not model alternative ways that individuals may finance their consumption in the face of an income shock. For example, we do not account for the scenario where an increase in the penalty rate for early withdrawals drives individuals with few other liquid savings to take on additional debt, perhaps at a high interest rate. Our results suggest that discouraging early withdrawals from retirement accounts through lower penalties that are in place for a longer period of time, and finding other policies to help individuals through unexpected financial hardship before the age of penalty-free withdrawals, may ultimately increase overall welfare.

7 Conclusion

We provide empirical evidence that traditional IRA holders respond to both the early withdrawal penalty and the excess accumulation penalty. We use the observed bunching response to estimate the magnitude of the response. We find that, every year, approximately 1.4% of traditional IRA holders (about 648,400 individuals) change the timing of withdrawals as a result of the early withdrawal penalty, and 16.2% of traditional IRA holders (about 7.7 million individuals) change the timing of withdrawals in response to the excess accumulation penalty. These shifts impact at least $57.9 billion in withdrawals each year. We believe these estimates should be considered lower bounds of the number of people whose behavior changes as a result of these penalties.

After estimating our model, we consider an array of counterfactual policies in which we alter the age threshold and penalty rate for the two penalties. We find that increasing the age for penalty-free withdrawals to be between $62\frac{1}{2}$ and $65\frac{1}{2}$, while lowering the penalty rate to 5%, increases both welfare and tax revenue as well as IRA balances at age $65\frac{1}{2}$. Similarly, we find that increasing the age for Required Minimum Distributions to be between $71\frac{1}{2}$ and $73\frac{1}{2}$, with a penalty rate of 50% or higher, increases both welfare and the present discounted value of lifetime income tax remittances.

Our results have three important implications for policy makers considering using these penalties as policy levers to increase retirement savings. First, there are combinations of age threshold and penalty rate for these penalties that increase both average welfare and tax remittances.

Second, encouraging individuals to keep their money in retirement savings accounts as long as possible increases welfare, as long as the penalty for early withdrawal isn't too high. Alternative strategies to assist individuals facing unexpected financial shocks before retirement should be considered in lieu of policies that make it easier for individuals to take early withdrawals from their retirement accounts.

Finally, the people who react to the early withdrawal penalty may differ from those who only take withdrawals because of Required Minimum Distributions. Intuitively, the early withdrawal penalty affects people who are trying to access their savings earlier in life, whereas the excess accumulation penalty affects individuals who want to enjoy the tax benefits for as long as possible. In addition, IRA holders have, on average, considerably higher incomes than non-IRA holders. If we want to create policy to help people save for retirement and we are worried that people with lower income might be more at risk for not having enough savings, then changing the policy levers associated with these accounts may not get us far enough. Incorporating extensive margin decisions into the model is a direction for future work.

References

Amromin, Gene and Paul Smith, "What Explains Early Withdrawals from Retirement Accounts? Evidence from a Panel of Taxpayers," *National Tax Journal*, 2003, *56* (3), 595–612.

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

Bee, Adam and Joshua Mitchell, "Do Older Americans Have More Income Than We Think?," U.S. Census Bureau Working Paper SEHSD-WP2017-39 2017.

Benartzi, Shlomo and Richard H. Thaler, "Myopic Loss Aversion and the Equity Premium Puzzle," *The Quarterly Journal of Economics*, 1995, *110* (1), 73–92.

Best, Michael Carlos, James S Cloyne, Ethan Ilzetzki, and Henrik J Kleven, "Estimating the Elasticity of Intertemporal Substitution Using Mortgage Notches," *The Review of Economic Studies*, 2019, *87* (2), 656–690.

Brady, Peter J. and Steven Bass, "Reconciling Form 1040 and Form 1099-R Data," SOI Working Paper 2020.

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

Butrica, Barbara A., Sheila R. Zedlewski, and Philip Issa, "Understanding Early Withdrawals from Retirement Accounts," The Urban Institute Retirement Policy Program Discussion Paper 10-02 2010.

Chetty, Raj, "A New Method of Estimating Risk Aversion," *American Economic Review*, December 2006, *96* (5), 1821–1834.

_ , John N. Friedman, Tore Olsen, and Luigi Pistaferri, "Adjustment Costs, Firm Response, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records," *Quarterly Journal of Economics*, 2011, *126*, 749–804.

Choukhmane, Taha, "Default Options and Retirement Saving Dynamics," Working Paper 2021.

Einav, Liran, Amy Finkelstein, and Paul Schrimpf, "The Response of Drug Expenditure to Nonlinear Contract Design: Evidence from Medicare Part D," *Quarterly Journal of Economics*, 2015, *130* (2), 841–899.

Feldstein, Martin, "Inflation, Income Taxes, and the Rate of Interest: A Theoretical Analysis," *The American Economic Review*, 1976, *66* (5), 809–820.

_ **, Jerry Green, and Eytan Sheshinski**, "Inflation and Taxes in a Growing Economy with Debt and Equity Finance," *Journal of Political Economy*, 1978, *86* (2), 809–820.

French, Eric, "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behavior," *The Review of Economic Studies*, 2005, *72* (2), 395–427.

Goda, Gopi Shah, Damon Jones, and Shanthi Ramnath, "How Do Distributions from Retirement Accounts Respond to Early Withdrawal Penalties? Evidence from Administrative Tax Returns," Working Paper 2018.

Goodman, Lucas, Kathleen Mackie, Jacob Mortenson, and Heidi Schram, "The Evolution of Leakage and Retirement Asset Flows in the U.S.," Working Paper 2019.

Gourinchas, Pierre-Olivier and Jonathan A. Parker, "Consumption Over the Life Cycle," *Econometrica*, 2002, *70* (1), 47–89.

Government Accountability Office, "Retirement Security: Most Households Approaching Retirement Have Low Savings," GAO-15-419 2015.

Guvenen, Fatih, "Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective," *Journal of Monetary Economics*, 2006, *53* (7), 1451 – 1472.

Internal Revenue Service, "2016 Individual Income Tax Returns Line Item Estimates," Technical Report 2016.

Investment Company Institute, "Investment Company Factbook," Technical Report 2019.

Jakobsen, Katrine, Kristian Jakobsen, Henrik Kleven, and Gabriel Zucman, "Wealth Taxation and Wealth Accumulation: Theory and Evidence From Denmark," *The Quarterly Journal of Economics*, 10 2019, *135* (1), 329–388.

Kleven, Henrik J. and Mazhar Waseem, "Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan," *Quarterly Journal of Economics*, 2013, *128*, 669–723.

Kleven, Henrik Jacobsen, "Bunching," *Annual Review of Economics*, 2016, *8* (1), 435–464.

Laibson, David I., Andrea Repetto, Jeremy Tobacman, Robert E. Hall, William G. Gale, and George A. Akerlof, "Self-Control and Saving for Retirement," *Brookings Papers on Economic Activity*, 1998, *1998* (1), 91–196.

Laibson, David, Peter Maxted, Andrea Repetto, and Jeremy Tobacman, "Estimating Discount Functions with Consumption Choices with the Lifecycle," Working Paper 2018.

Larrimore, Jeff, Jacob Mortenson, and David Splinter, "Household Incomes in Tax Data: Using Addresses to Move from Tax Unit to Household Income Distributions," *Journal of Human Resources*, 2019.

Mehra, Rajnish and Edward C. Prescott, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 1985, *15* (2), 145–161.

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

OASDI, "The 2020 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds," Available at https://www.ssa.gov/oact/tr/2020/. 2020.

Rouwenhorst, Geert K., "Asset Pricing Implications of Equilibrium Business Cycle Models," in Thomas Cooley, ed., *Structural Models of Wage and Employment Dynamics*, Princeton: Princeton University Press, 1995.

Sabelhaus, John, "Modeling IRA Accumulation and Withdrawals," *National Tax Journal*, 2000, *53* (4), 865–876.

TIGTA, "Improvements Can Be Made to Educate and Notify Taxpayers of Required Minimum Distributiona Requirements From Individual Retirement Arrangements," Technical Report 2015-10-042.

A Additional institutional information about IRAs

(For online publication)

A.1 Intuition for IRA tax benefit

Consider an individual that wants to invest an amount a of taxable income. She faces an average income tax rate of t. She can either invest the amount in a regular savings account after remitting income tax, or in a tax-benefited savings account before remitting income tax. Both accounts face the same annual rate of return r. The growth of these accounts is outlined in Table 9:

Table 9: Account growth over time in simple example

Period	Tax-benefited account		Regular account: No tax on return		Regular account: With tax on return
0	a		$a(1-t_i)$		$a(1-t_i)$
1	$a(1+r)$		$a(1-t_i)(1+r)$		$a(1-t_i)(1+r-rt_r)$
2	$a(1+r)^2$		$a(1-t_i)(1+r)^2$		$a(1-t_i)(1+r-rt_r)^2$
3	$a(1+r)^3$		$a(1-t_i)(1+r)^3$		$a(1-t_i)(1+r-rt_r)^3$
...					
Account balance in period before withdrawing					
$k-1$	$a(1+r)^{k-1}$		$a(1-t_i)(1+r)^{k-1}$		$a(1-t_i)(1+r-rt_r)^{k-1}$
Amount after withdrawal					
k	$a(1-t_i)(1+r)^k$	$=$	$a(1-t_i)(1+r)^k$	$>$	$a(1-t_i)(1+r-rt_r)^k$

If the individual withdraws from the tax-benefited account, she will owe income tax and her take-home value will be equal to the balance in the regular account. However, this amount is not taxed further. In contrast, there are two additional taxes the individual faces if she holds her money in the regular account depending on how the money is invested. She may owe tax on the interest accumulated each year she held the account. Upon withdrawal, she may owe capital gains taxes. Neither of these apply to the tax-benefited account. As a result, it is as if the returns to the tax-benefited account were not subject to additional tax beyond income tax.

A.2 Roth IRAs

The primary difference between traditional IRAs and Roth IRAs is the timing of when income tax is due. Contributions to Roth IRAs are made after income tax is remitted, with no tax due on qualified distributions (including on any earnings between contribution and withdrawal). Roth IRAs are not subject to Required Minimum Distributions (withdrawals are not required from Roth IRAs until after the death of the account owner). Roth IRA holders do face the early withdrawal penalty, but the penalty only applies on withdrawals that are larger than the principal investment.

There are a few other differences between traditional IRAs and Roth IRAs. All taxpayers who earn sufficient income to make a contribution to a traditional IRA are eligible to hold a traditional IRA, whereas eligibility to contribute to a Roth IRA is phased out for taxpayers with high enough income (see Appendix A.4). There are also differences in when and what taxes are owed upon death of the account holder (see Appendix A.6).

A.3 SEP and SIMPLE IRAs

There are two types of IRAs in addition to traditional IRAs and Roth IRAs: SIMPLE IRAs and SEP IRAs. SIMPLE IRAs allow both employee and employer to contribute to a traditional IRA set up for the employee, and are intended to be used by small employers (<100 employees) not sponsoring a retirement plan such as a 401(k). SEP (Simplified Employee Pension) IRAs are used primarily by self-employed and small business owners. Only the employer contributes, and the employee is always 100% vested. If an employer offers SEP IRAs to their employees, the employer must contribute for all employees that meet a set of requirements.[39]

Table 10 shows the number of taxpayers that made contributions to or took distributions from each of these account types, as well as the total end-of-year fair market value of these accounts, in 2015. Because these IRAs make up a small fraction of total IRA accounts and face the same penalty structure as traditional IRA accounts, we follow Mortenson et al. (2019) and count contributions toward (and distributions from) a SEP or SIMPLE IRA as contributions toward (and distributions from) a traditional IRA for our analysis.

A.4 Contribution and deductibility limits

Contributions into IRAs are capped. In 2015, individual taxpayers could not contribute more than $5,500 in total to traditional and Roth IRA accounts. That amount was higher for individuals over age 50 ($6,500). Contributions to IRAs are limited to an individual's taxable compensation, so if that was less than $5,500 for an individual, that individual's contribution limit would be equal to their taxable compensation. Taxpayers were not allowed to make contributions to a traditional IRA after age $70\frac{1}{2}$ in 2015; contributions to a Roth IRA were allowed after age $70\frac{1}{2}$. In the event a taxpayer makes a contribution

[39] A third additional type of IRA–the SARSEP (Savings Incentive Match Plan for Employees) IRA–are SEP IRAs established before 1997 that include a salary reduction arrangement.

Table 10: IRA activity by type of account, 2015

IRA type	Contributions		Distributions		End-of-year FMV	
	Taxpayers	Pct.	Taxpayers	Pct.	FMV $1,000	Pct.
Traditional	4,305,106	31.6%	17,360,396	92.4%	6,386,720	85.4%
Roth	6,363,335	46.7%	708,221	3.8%	625,077	8.4%
SEP	1,093,512	8.0%	670,990	3.6%	364,264	4.9%
SIMPLE	1,865,777	13.7%	40,459	0.2%	101,194	1.4%
Total	13,627,730	100.0%	18,780,066	100.0%	7,477,255	100.0%
Total (unique)	13,006,314	n/a	18,670,599	n/a	n/a	n/a

Source: Authors' calculations from Statistics of Incomes Tax Stats: Accumulation and Distribution of Individual Retirement Arrangements (IRA) Table 1 (2015), available at https://www.irs.gov/pub/irs-soi/15in01ira.xls. Total (unique) count is less than the total because some taxpayers may have made contributions or distributions from more than one account. Contributions do not include rollovers.

above the limit (or any contribution at age $70\frac{1}{2}$ or older), the excess amount is taxed at 6% for each year it remains in the IRA.

When money in one retirement account is moved into a different retirement account, it is called a "rollover." For example, if you leave job A, you can transfer the amount in your 401(k) from job A into the 401(k) you set up with your new job (or into your own IRA) and it would not count toward your contribution limit. Rollovers from, e.g., a 401(k) into an IRA are not counted toward the taxpayer's contribution limit.

Contributions to Roth IRAs are not allowed for taxpayers above certain income levels. Individual taxpayers with a modified adjusted gross income (MAGI) of less than $116,000 in 2015 were allowed to contribute the full $5,000 to a Roth IRA. Between a modified AGI of $116,000 and $131,000, taxpayers were eligible to contribute a reduced amount of their contribution limit to a Roth IRA. The partial amount for an individual taxpayer in 2015 can be calculated as:

$$\text{partial contribution} = \text{contribution limit} - \left(\frac{\text{MAGI} - \$116,000}{\$10,000} \cdot \text{contribution limit}\right).$$

After a modified AGI of $131,000, taxpayers wishing to contribute to an IRA were forced to contribute to a traditional IRA. For married taxpayers filing jointly, the upper bound of modified AGI for contributing the full amount to a Roth IRA was $183,000, and the lower bound after which IRA contributions must be made to a traditional account was $193,000. One strategy expected to be used by some taxpayers who earn too much to contribute to a Roth IRA is to contribute to a traditional IRA and convert it to a Roth IRA later. No taxes are assessed with a rollover, but if you rollover from a traditional account into a

Roth account, you must include the distributed amount as income (i.e., it will be included in your taxable income that year).

The deductibility of contributions to a traditional IRA are limited for taxpayers who are also covered by an employer-sponsored plan, such as 401(k). In 2015, only individual taxpayers with a modified AGI of less than $61,000 ($98,000 for taxpayers filing jointly) were able to deduct the full contribution limit of $5,000 if the entire amount was contributed to a traditional IRA. The deductible amount phased out until a modified AGI of $71,000 ($188,000 for taxpayers filing jointly), at which point individual taxpayers covered by an employer-sponsored plan were not eligible to deduct any amount of a contribution to a traditional IRA.

A.5 Schedule for Required Minimum Distributions

After age $70\frac{1}{2}$, holders of traditional, SEP, and SIMPLE IRAs are required to take "minimum distributions" from their account. The first RMD must be taken by April 1 of the year after the calendar year in which the account holder turns $70\frac{1}{2}$; after that, RMDs are due by December 31 of each calendar year. This means that account holders who wait until the calendar year after the year in which they turn $70\frac{1}{2}$ to remit their first RMD will owe two RMD payments that year: one by April 1 and one by December 31.

The RMD amount for a given year is equal to the account balance on December 31 of the preceding year divided by a value known as a "distribution period." These values are given by the IRS's "Uniform Lifetime Table," given below in Table 11.[40] The estimated RMD for an IRA is reported on IRS Form 5498. If a taxpayer holds more than one IRA subject to RMDs, their total RMD is equal to the sum of the RMDs from each account. In this scenario, the taxpayer may take the total RMD from one account.[41]

A.6 Penalties and the death of the account holder

What happens to a traditional IRA after the account owner dies depends on who inherits the account and the age of the original account owner upon death. If a spouse inherits the account, the spouse is eligible to designate herself as the IRA account owner or keep the account in the decedent's name. In the former scenario, the surviving spouse is allowed to make contributions, take distributions, rollover the assets into a different account, and

[40] The distribution schedule is different if the account owner's spouse is the sole beneficiary of the account and the spouse is more than 10 years younger than the account holder. In this case, the distribution period value is determined by the account holder's age and the spouse's age on their respective birthdays in that calendar year.

[41] The same is not true for 401(k)s: RMDs from 401(k)s must be taken from each account separately.

Table 11: Uniform Lifetime Table from IRA Publication 590

Age	Distribution period	Age	Distribution period	Age	Distribution period	Age	Distribution period
70	27.4	82	17.1	94	9.1	106	4.2
71	26.5	83	16.3	95	8.6	107	3.9
72	25.6	84	15.5	96	8.1	108	3.7
73	24.7	85	14.8	97	7.6	109	3.4
74	23.8	86	14.1	98	7.1	110	3.1
75	22.9	87	13.4	99	6.7	111	2.9
76	22.0	88	12.7	100	6.3	112	2.6
77	21.2	89	12.0	101	5.9	113	2.4
78	20.3	90	11.4	102	5.5	114	2.1
79	19.5	91	10.8	103	5.2	115+	1.9
80	18.7	92	10.2	104	4.9		
81	17.9	93	9.6	105	4.5		

Source: IRA Required Minimum Distribution Worksheet, https://www.irs.gov/pub/irs-tege/uniform_rmd_wksht.pdf. This schedule has been in place since 2002. The schedule from 1999-2001 required slightly higher withdrawals than the schedule presented here.

even name a new beneficiary for the account. The surviving spouse is also subject to RMD rules based on their own age. If the surviving spouse keeps the account in the decedent's name, he or she must start taking distributions when the descendent would have turned $70\frac{1}{2}$ (or continue taking distributions if the decedent had already started taking RMDs), but based on his or her own life expectancy.

Non-spouse beneficiaries of traditional IRAs do not owe taxes on the assets in the account until they begin taking distributions. If the original account owner died after RMDs began, the beneficiary must continue to take distributions. The amount of the distribution is based off of the longer of the original account owner's and the beneficiary's life calculated expectancy. If the owner died before they were subject to RMDs, the beneficiary must either take RMDs based off of their own life expectancy starting the year after the original account owner's death, or elect to fully distribute the account under the "5-year rule."[42] The 5-year rule requires the IRA beneficiaries to withdraw all of the assets of the IRA by December 31 of the calendar year containing the fifth anniversary of the original account owner's death, but does not require any distributions before that time. Failing to fully distribute the assets after electing to distribute according to the 5-year rule results in a 50% excise tax on the remaining amount. The 5-year rule always applies

[42]If there are multiple non-spouse beneficiaries, the required distribution amounts are based on the age of the oldest beneficiary. If the beneficiaries split the original account into separate accounts, each beneficiary will be subject to RMD rules based on their own age.

in cases where the beneficiary isn't an individual (i.e., is a trust or an organization).

When the holder of a Roth IRA dies and the beneficiary is the spouse, the spouse becomes the new owner of the account. The spouse may choose not to take ownership of the account, in which case he or she faces the same options as a non-spouse beneficiary. A non-spouse beneficiary may either distribute the entirety of the Roth IRA by the end of the calendar year containing the fifth anniversary of the account holder's death, or over his or her own life expectancy. A surviving spouse is allowed to delay distributions until the deceased spouse would have turned $70\frac{1}{2}$. The ability of a non-spouse beneficiary to take distributions based on their life expectancy was curtailed by the SECURE Act, which was signed into law in December 2019. The SECURE Act includes a provision that, going forward, requires non-spouse beneficiaries to distribute the funds of inherited IRAs within 10 years. This is meant to prevent "Stretch IRAs," in which IRA accounts are passed down, e.g. to grandchildren, who are then able to take distributions over their expected lifespan and enjoy tax-free growth along the way.

B Data description

(For online publication)

This appendix serves as a guide to the data and terminology used in this project. Appendix B.1 defines relevant terms and, when applicable, provides additional information about how key variables were constructed. Appendix B.2 discusses our sample construction and the IRS forms from which we extracted the data.

B.1 Glossary of terms and variable construction

Contribution We get contribution information from Form 5498. We assume that anything labeled as a contribution that was larger than 110% of the taxpayer's contribution limit in that year was a data error. In this case, we coded the contribution as a rollover.

Defined-benefit (DB) plan Employer-sponsored savings plan that pays pre-determined benefits from a collective trust fund rather than an individualized account. Commonly referred to as pensions.

Defined-contribution (DC) plan Employer-sponsored savings plan with an individual account for each participant. When an individual retires, the designated account is used to provide benefits to that individual.

Distribution The word used by the IRS and financial industry to talk about withdrawing money from a tax-deferred retirement plan. We calculate distributions from IRA accounts using data from Form 1099-R. We pull all Form 1099-Rs associated with our random sample. We follow a standard algorithm for determining which Form 1099-Rs indicated distributions from traditional, Roth, and SIMPLE IRAs, and which Form 1099-Rs indicated rollovers from Roth accounts into traditional accounts. We exclude distributions from accounts other than IRAs, as well as the following types of IRA distributions:

- IRA distributions due to death of the account holder
- IRA distributions due to disability
- Recharacterizations of this or prior year's IRA distribution
- Corrective distributions due to IRA contributions or deferrals above the annual limit

These exclusions eliminate 5.7% from our sample of IRA-related Form 1099-Rs, the majority of which come from distributions due to death of the account owner (4.8%).

For each taxpayer, we determine the taxable distributions and gross distributions for each type of IRA, as well as gross distributions excluding any distributions not subject to the early withdrawal penalty. Out of 22,192,136 Form 1099-Rs associated with our sample, there are 388 cases where the taxable amount is greater than the gross amount. In these cases, we assume the taxable amount is correct and replace the gross amount with the taxable amount, unless the difference between the two is greater than $1,000$ and the value of the gross amount is greater than 500. When this occurs, we add the value of the taxable amount to the gross amount.

Early withdrawal penalty The penalty owed to the IRS when the owner of an individual retirement account takes a distribution before the age of $59\frac{1}{2}$. The penalty is equal to 10% of the amount withdrawn. Taxpayers may receive an exemption if they take an early withdrawal from an IRA to pay for qualified medical or higher education expenses. First-time homebuyers may take a distribution of up to $10,000 from an IRA without penalty.

Employer-sponsored plan (ESP) Can be either defined-contribution (DC) or defined-benefit (DB). Figure 11 presents evidence that DC employer-sponsored accounts such as 401(k)s are increasingly popular over DB accounts such as pensions. In each of the panels, DB plans are shown by the solid line and DC plans by the dashed line. Figure 11a shows the total number of plans offered (in thousands), Figure 11b shows the total number of active participants (in millions), and 11c shows the total assets in these plans (in trillions). For all three measures, it's clear that DC plans are growing faster than DB plans (and, in some cases, that DB plans are shrinking).

Figure 11: Trends in employer-sponsored DB vs. DC plans

(a) Number of plans (b) Active participants (c) Total assets

Source: Tables E1, E7, and E10 of *Private Pension Plan Bulletin Historical Tables and Graphs 1975-2017*, respectively, published by the U.S. Department of Labor. Defined-contribution plans are shown by the dashed line; defined-benefit plans are shown by the solid line. In 2004, the definition of who was participating was revised.

Excess accumulation penalty The penalty owed to the IRS when the owner or beneficiary of an individual retirement account fails to take a required minimum distribution. The penalty is equal to 50% of the required amount not withdrawn. The penalty may be waived when the excess accumulation was the result of a reasonable error and the account holder took steps to correct the insufficient distribution.

Fair-market value (FMV) The value of an individual retirement account if all of the assets were to be sold as reported on IRS Form 5498. For our purposes, equivalent to the account balance.

Individual Retirement Account (IRA) Also known as an Individual Retirement Arrangement. IRAs are tax-benefited retirement savings accounts. We define anyone in our sample that we observe making a contribution to an IRA (including rollovers), taking a distribution from an IRA, or having a positive IRA balance as an "IRA account holder."

Lump-sum distribution The distribution of the entire balance of a taxpayer's ESP(s) from a single employer.

Qualified charitable distribution (QCD) Direct transfer from an IRA account to a qualified charity. These distributions can be counted toward RMDs. Conditions: must be $70\frac{1}{2}$ or older at the date of distribution, total amount contributed not to exceed the amount that would otherwise be taxed as ordinary income up to the maximum amount ($100,000), which applies to the sum of all QCDs made. A distribution made to an account holder that is then donated to charity does not count as a QCD.

Required Minimum Distribution (RMD) Starting in the tax year that an account holder turns $70\frac{1}{2}$, holders of traditional 401(k) and IRA accounts are required to make a minimum distribution every year. Late RMD payments are taxed at 50%. Amount is based on account balance on Dec. 31 of previous tax year and life expectancy tables. First payment is due by April 1 of the year after you turn $70\frac{1}{2}$; subsequent payments must be made by Dec. 31 of each calendar year. As a result, taxpayers may have two RMDs included in their taxable income in the first calendar year they make RMD payments. In December 2019, the Setting Every Community Up for Retirement Enhancement (SECURE) Act was signed into law. Section 114 increases the age at which taxpayers are subject to excess accumulation penalties for failing to take required distributions from $70\frac{1}{2}$ to 72. Because we consider a time period before 2019, we use $70\frac{1}{2}$ as the age threshold throughout the paper.

Rollover The word used when money in one retirement account is moved into a different retirement account. For example, if you leave job A, you can transfer the amount in your 401(k) from job A into the 401(k) you set up with your new job (or into your own IRA). No taxes are assessed with a rollover unless it is a Roth conversion. Rollovers are not counted toward contribution limits.

When a taxpayer rolls over funds from, e.g., a 401(k) or another IRA into an IRA, a Form 1099-R will be issued to mark the distribution. If the rollover is a Roth conversion, the distributions will be marked as a rollover on Form 1099-R and a Roth conversion on Form 5498. The Form 5498 is issued to mark the contribution to the IRA.

Roth conversion A type of rollover when funds in a traditional IRA are rolled over into a Roth IRA. For this type of rollover, income tax is due on the distribution (i.e., you must include the distributed amount as income (i.e., it will be included in your taxable income that year).

Roth IRA A type of IRA in which income tax is remitted before funds are moved into the account. Income taxes are not due on qualified distributions, including on returns to the principal.

Savings (initial distribution of non-tax-benefited) IRS Form 1099-INT is used to report interest income to the IRS. We add together the values of Box 1 ("Interest income") and Box 3 ("Interest on U.S. Savings Bonds and Treasury obligations"), both winsorized at the 1^{st} and 99^{th} percentiles, as reported on IRS Form 1099-INT for the individuals in our sample at age 40. We estimate assets in regular savings by dividing that amount by

0.01. The average estimated value of regular savings using this method is \$57,539. This is well within the ballpark of the average holdings in transaction accounts, Certificates of Deposit, and savings bonds estimated by the Federal Reserve, which in 2016 was estimated to be \$64,800 for families with a head of household between 35 and 45. Source: authors' calculations using the U.S. Federal Reserve Board, Historic Tables and Charts - Tables Based on Internal Data - Estimates Inflation Adjusted to 2016 Dollars, Table 6, available at https://www.federalreserve.gov/econres/scfindex.htm.

Total distribution When all of the contents of an individual retirement account are withdrawn.

Traditional IRA A type of IRA in which funds are contributed before income tax is remitted. Income tax is then due on all qualified distributions, including on returns to the principal.

Withdrawal See distribution.

B.2 Sample construction

We take advantage of administrative tax data in this project. Our initial panel is a 5% random sample of individuals aged 18 or older in 1999. We follow these taxpayers through 2015. The 17-year panel is balanced apart from exit due to death and emigration.

Our initial sample includes 14,606,095 individual taxpayers. We make several cuts to this initial sample. We exclude taxpayers who died before 2000 and therefore are never alive during our sample period. We also exclude taxpayers who were older than 90 (born before 1909) in 1999. We exclude any taxpayer whose taxpayer identification number (TIN) was ever associated with an invalid TIN flag or multiple Form 1040 filings.[43] These last two exclusions in particular are to deal with issues of identity theft. Table 12 summarizes the impact these exclusions have to our sample size. Our final sample includes 11,950,600 individual taxpayers, 32.7% of which are IRA account holders. We identify taxpayers as "IRA account holders" if we ever observe them making a contribution to an IRA (including a rollover), taking a distribution from an IRA, or having an outstanding IRA balance reported on Form 5498.

We focus on IRAs because we can cleanly identify IRA contributions and distributions in the administrative tax data. Form 5498 is intended for contributions to IRAs only, and separately identifies contributions to traditional and Roth IRAs. Form 5498 is filed by

[43]We understand that this necessarily excludes individuals choosing to file Married Filing Separately.

Table 12: Data exclusions

Exclusion	Unique taxpayers	Percent of initial sample
Original sample	14,606,095	100.0%
Died before 2000	456,240	3.1%
Born before 1909	1,072,032	7.3%
Invalid TIN	1,126,383	7.7%
Multiple F-1040s	840	0.0%
Final sample	11,950,600	81.8%
IRA holders	3,913,401	
		Percent of IRA sample
Traditional IRA holders	2,790,313	72.2%
Roth IRA holders	1,071,964	27.8%
Both types	616,020	15.7%

the individual or organization who made contributions on behalf of the recipient. For example, the bank that manages Joe's IRA would submit Form 5498 for Joe if Joe made a contribution to his IRA. In addition to contribution information, Form 5498 includes the end-of-year fair-market values (FMV) of an IRA. Financial institutions are required to submit Form 5498 each year for an existing IRA even if no contribution is made. There is some debate about whether or not financial institutions actually do this, but our data suggest that most financial institutions do submit Form 5498 even if the account holder did not make a contribution that year.

Line 7 of Form 1099-R includes a checkbox to indicate that the distribution comes from an IRA. We use the codes entered in Line 7 to then distinguish between distributions from traditional and Roth IRAs. Notably, we can separately identify normal distributions, early distributions without a known exception, and early distributions with a known exception. We are also able to identify rollovers to traditional and Roth IRAs. In contrast, neither Form 1099-R nor Form 1040 separately identify distributions from pensions and annuities from distributions from 401(k)s and other employer-sponsored defined-contribution accounts. Form 1099-R is filed by the institution who issued the distribution. If Joe took a distribution from his IRA, the bank that manages the IRA would be responsible for submitting Form 1099-R to the IRS.

IRA distributions are also reported on Lines 15a and 15b of Form 1040. We prefer to use the amounts from the third-party reported Form 1099-R. Using Form 1099-R has two primary advantages. First, we are able to capture distributions from non-filers. Second, Lines 15a and 15b on Form 1040 aggregates distribution amounts for spouses for married taxpayers filing jointly. We observe that, for single taxpayers, the total amount distributed from traditional IRAs as reported on Form 1099-R and the taxable amount reported on Line 15b of Form 1040 are within 1% of each other in 78% of cases, and within $100 of each other in 94% of cases. This is in line with what is reported in Brady and Bass (2020).

We also use data from individual income tax returns. In particular, we use data from Forms 1040 and 5329. Form 1040 includes information such as total wages and adjusted gross income. Schedules A and C for Form 1040 include charitable giving (for itemizers) and self-employment income, respectively. Form 5329 is used to declare additional taxes owed due to early withdrawal, non-qualified distributions, or failing to receive a minimum required distribution. Finally, we use data from Form 1099-INT to estimate the initial distribution of non-IRA savings in the model.

C Additional evidence of bunching

(For online publication)

C.1 Diagnostic tests

We conduct two diagnostic tests to confirm that our observed bunching is a result of the age thresholds for penalties related to IRA withdrawals. In Appendix C.1.1, we plot the proportion of taxpayers receiving two placebo sources of income (wages and Social Security) at each age in 2005. In Appendix C.1.2, we separately plot the proportion of taxpayers taking a distribution from an IRA for taxpayers whose half-birthday is in the first or second half of the year. The results of both of these diagnostic tests suggest that taxpayers are changing their behavior in response to the penalties at ages $59\frac{1}{2}$ and $70\frac{1}{2}$ (i.e., the bunching we presented in Section 3 was not a result of other policies related to timing of income).

C.1.1 Placebo outcomes

We create the same figure as Figure 3a for two placebo outcomes: Social Security payments and wages. Figure 12a shows the proportion of taxpayers in our sample of IRA account holders who receive a social security payment at that age. Similarly, Figure 12b shows the proportion of taxpayers in our sample of IRA account holders who receive a wage at that age. In both cases, there is no bunching the proportion of taxpayers at either $59\frac{1}{2}$ or $70\frac{1}{2}$. This implies that the jumps we are seeing at $59\frac{1}{2}$ and $70\frac{1}{2}$ are not confounded with other changes in financial circumstances during retirement or pre-retirement.

Figure 12: Fraction of traditional IRA holders receiving Social Security or wages in 2005, by age

(a) Social Security benefits (b) Wages

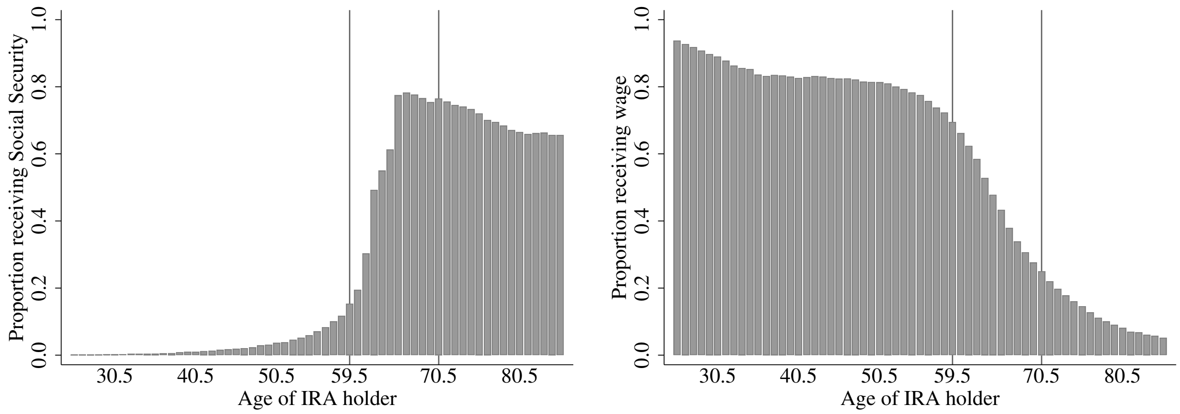

Notes: N = 1,449,868 taxpayers. Limited to taxpayers with a positive traditional IRA balance in 2004 or 2005.

C.1.2 Timing of birthday

In our second diagnostic, we take advantage of the fact that the age thresholds kick in at taxpayers' half birthdays. We compare traditional IRA holders who turn age X and X.5 in the same calendar year versus traditional IRA holders who turn age X and X.5 in different calendar years. Individuals who turn 59 and $59\frac{1}{2}$ in the same year are eligible to take withdrawals penalty-free in that year, whereas individuals who turn 59 and $59\frac{1}{2}$ in different years are not eligible to take distributions penalty-free in the year they turn 59. If taxpayers are delaying withdrawals as a result of the age threshold, we should see a spike in the proportion of taxpayers withdrawing at age 59 for taxpayers that turned 59 and $59\frac{1}{2}$ in the same calendar year, but at age 60 for taxpayers that turned $59\frac{1}{2}$ in the calendar year after the year in which they turned 59.

Figure 13 shows that this is exactly what happens. Figure 13a shows taxpayers whose half-birthday is in the same calendar year as their birthday. For these taxpayers, we see increases in withdrawal rates at ages 59 and 70. Figure 13b shows taxpayers whose half-birthday is in a different calendar year than their birthday. For these taxpayers, we see increases in withdrawal rates at ages 60 and 71. This is consistent with taxpayers changing their behavior as a result of exactly when these specific age thresholds apply.

Figure 13: Fraction of traditional IRA holders taking a distribution in 2005, by timing of half-birthday

(a) Half-birthday in same calendar year (b) Half-birthday in next calendar year

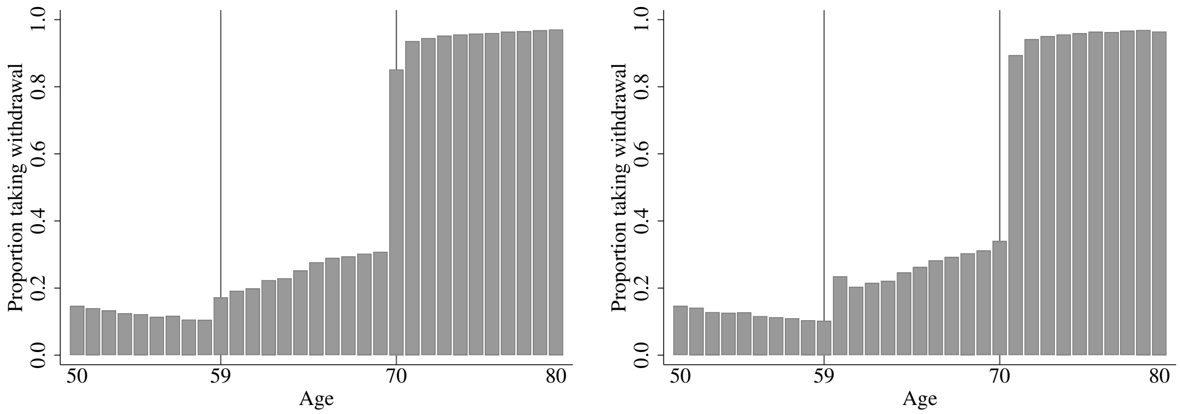

Notes: Figure 13a: N = 490,904 unique taxpayers. Figure 13b: N = 517,808 unique taxpayers. Limited to taxpayers with a positive traditional IRA balance in 2004 or 2005.

C.2 Roth IRAs

We created the same figures for Roth IRAs as we do to demonstrate bunching for traditional IRAs. We do not observe the same clear bunching for Roth IRAs as we do for traditional IRAs.

Figure 14 shows the same empirical distribution as Figure 1 in Section 3.2, but for Roth IRA account holders. Figure 14a shows the number of Roth IRA holders taking their first distribution from a Roth IRA at each age; Figure 14b shows total amount taken at those distributions.

There are a few key things to note about this figure. First, the scale of these graphs is considerably different than that for traditional IRAs. Figure 14a goes through 8 thousand, whereas Figure 1a is scaled through 150 thousand. Similarly, Figure 14b is scaled through $50 million, whereas Figure 1b is scaled through $1.8 billion. We also observe small spikes at age $59\frac{1}{2}$, but not at $70\frac{1}{2}$. This is consistent with the fact that the early withdrawal penalty applies to Roth IRAs, but the excess accumulation penalty does not.

Figure 15 shows the same empirical distribution as Figure 3 in Section 3.3, but for Roth IRA holders. The first panel shows the proportion of Roth IRA account holders taking a distribution in 2005; the second panel shows the average value of those distributions.

There are two things to note about these figures. First, we do not observe the same striking bunching at age $59\frac{1}{2}$ that we observe for traditional accounts. Second, there is an unexpected uptick in the number of Roth account holders who take their first distribution

Figure 14: Age of first distribution from a Roth IRA

(a) Number of Roth IRA holders (b) Gross distribution amount ($milion)

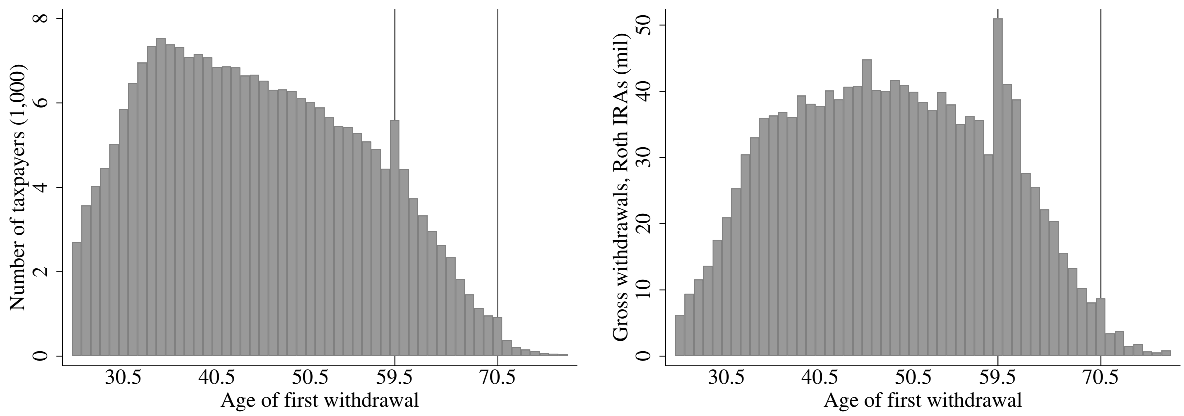

Notes: N = 236,062 unique taxpayers. Excludes taxpayers older than age 65 at the beginning of our sample period, and all data from 2009. Figure 14a shows the number of taxpayers in our sample we observe taking their first distribution from a Roth IRA at each age, excluding early distributions (before age $59\frac{1}{2}$) with a known, qualifying exception. Figure 14b shows the total amount withdrawn in the first distribution by Roth IRA account holders who took their first withdrawal at each age. Distribution amounts are adjusted for 2015 dollar values. Because taxpayers may hold both types of accounts, some taxpayers may appear in both Figure 1 (showing traditional IRAs in the main text) and this figure.

Figure 15: Distribution behavior in 2005: Roth IRAs

(a) Proportion of Roth IRA holders (b) Average distribution amount

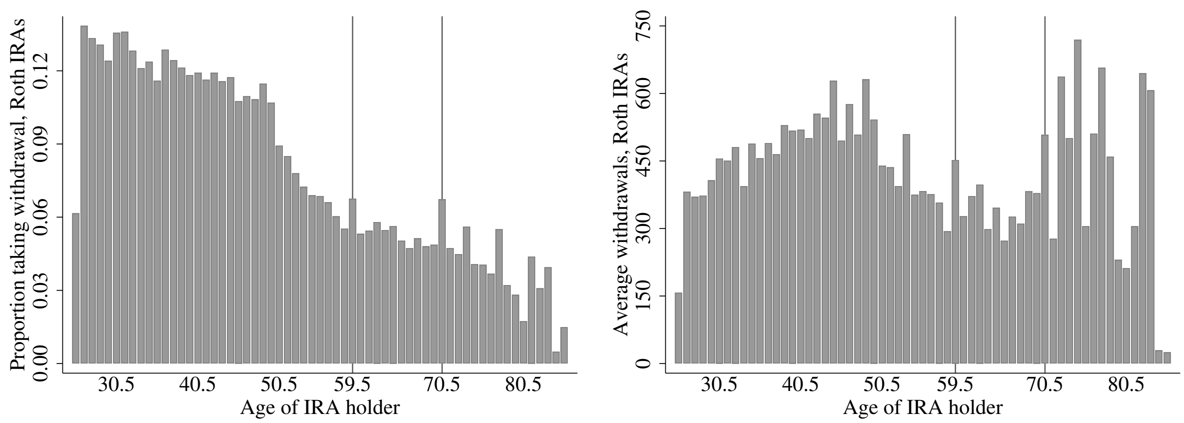

Notes: *Notes*: N = 269,353 unique taxpayers. Excludes early distributions (before age $59\frac{1}{2}$) with a known, qualifying exception. Excludes taxpayers older than age 65 at the beginning of our sample period. Limited to taxpayers with a positive Roth IRA balance in 2004 or 2005. Figure 15a shows the proportion of Roth IRA account holders taking a distribution in 2005, while Figure 15b show the total amount distributed from Roth IRA by age in 2005. Distribution amounts are inflated to 2015 values. Because taxpayers may hold both types of accounts, some taxpayers may appear in both Figure 3 (showing traditional IRAs in the main text) and this figure.

from a Roth IRA at age $70\frac{1}{2}$ (14a), and in the proportion of Roth account holders taking a distribution from a Roth IRA at age $70\frac{1}{2}$ (15a). This observation is unexpected because Roth IRA holders are not bound to the RMD rules. Figure 16 splits Roth IRA account holders into those that also hold a traditional IRA (solid line) and those that only hold a Roth IRA (dashed line). If it's true that IRA account holders with both types of accounts are more likely to take non-rollover distributions from a Roth account at age $70\frac{1}{2}$ than IRA account holders who only have a Roth account, it would suggests that these taxpayers may not realize the Required Minimum Distribution rules do not apply to Roth IRAs.

Figure 16a shows the number of taxpayers taking their first distribution by age. We see that there is an increase in the number of taxpayers taking their first distribution at age $70\frac{1}{2}$ for Roth IRA account holders that also have a traditional IRA, but not for Roth account holders who do not have a traditional IRA. Figure 16b shows the proportion of Roth IRA holders taking a distribution at each age. While we observe an uptick at $70\frac{1}{2}$ for both groups, the trend for Roth IRA account holders who do not also have a traditional IRA is generally more erratic than that of Roth IRA account holders who do show evidence of also having a traditional IRA. These figures suggest that it may be true that taxpayers that hold both types of accounts are more likely to take distributions from their Roth IRA account at age $70\frac{1}{2}$ even though they are not required to.

Figure 16: Distribution behavior of Roth IRA account holders by whether or not they also hold a traditional IRA

(a) Age of first distribution from Roth IRA

(b) Proportion taking distributions

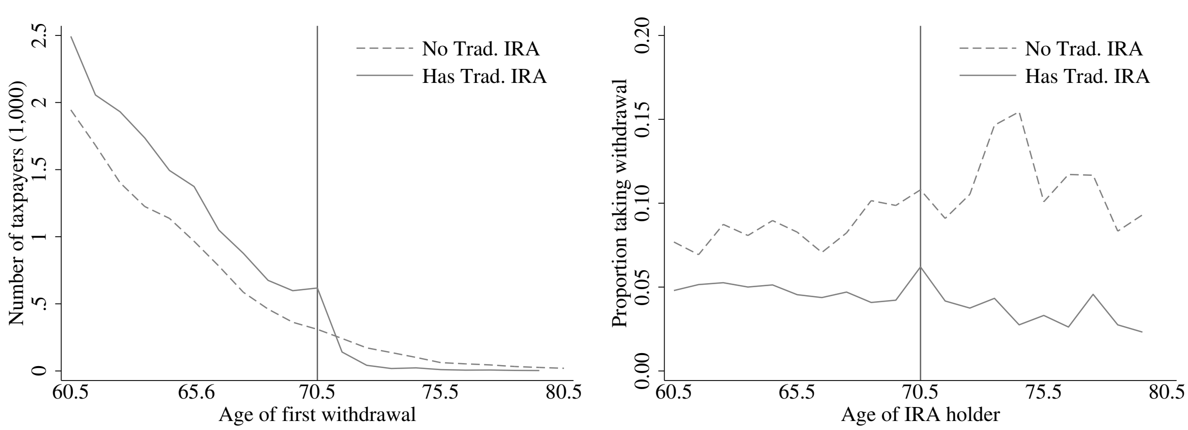

Notes: Figure 16a: N = 26,904 taxpayers; excludes taxpayers older than age 65 at the beginning of our sample period. Figure 16b: N = 50,756; limited to taxpayers with a positive Roth IRA balance in 2004 or 2005. The dotted line represents Roth IRA account holders who do not also hold a traditional IRA. The solid line represents Roth IRA account holders who do show evidence of holding a traditional IRA.

D Comparing different values of $\mathbf{a_{59.5,-}}$ and $\mathbf{a_{59.5,+}}$

(For online publication)

We estimated $\widehat{B}_{59.5}$ using the values of $a_{59.5,-}$ and $a_{59.5,-}$ that minimized the difference between $\widehat{B}_{59.5}$ and $\widehat{M}_{59.5}$. We find that, for the number of taxpayers impacted, $a_{59.5,-} = 54.5$ and $a_{59.5,+} = 61.5$. Table 13 shows the absolute value of $\widehat{B}_{59.5} - \widehat{M}_{59.5}$ for all possible combinations of $a_{59.5,-}$ and $a_{59.5,+}$ for the number of taxpayers that change when they take their first distribution. Table 14 shows the same calculation, but for amount of money taken out at those first distributions. Similarly, Table 15 shows the absolute value of $\widehat{B}_{59.5} - \widehat{M}_{59.5}$ for all possible combinations of $a_{59.5,-}$ and $a_{59.5,+}$ for the proportion of taxpayers taking a withdrawal by age in 2005, while Table 16 shows the same calculation for the average withdrawal amount by age in 2005.

Table 13: Robustness check: $\widehat{B}_{59.5} - \widehat{M}_{59.5}$ for reduced-form analysis: number of taxpayers, first distribution (1000s)

	$a_{59.5,+}$										
$a_{59.5,-}$	59.5	60.5	61.5	62.5	63.5	64.5	65.5	66.5	67.5	68.5	69.5
25.5	448.9	797.5	773.3	691.4	603692.6	2163406.3	2038868.1	270894.3	1302.2	1325.4	9631493.1
26.5	9794.1	521.8	12571.6	44365.1	153170.0	312354.3	226137.8	65355.0	1283.0	1295.2	350244.7
27.5	987.1	99.4	1174.7	1924.7	4593.9	5207.5	1044.2	1221.9	1269.2	1268.5	162897.1
28.5	1002.2	222.8	1166.9	1658.5	3018.9	2761.9	642.0	961.7	917.0	28.7	9216.3
29.5	1037.6	579.1	1158.8	1365.9	1816.9	1590.2	197.0	274.9	22.1	1363.8	7727.6
30.5	1115.8	898.7	1214.0	1315.9	1527.6	1392.0	585.9	301.6	511.3	1689.6	7380.5
31.5	1133.5	1013.1	1205.5	1267.6	1405.3	1301.0	668.3	436.0	627.1	1717.0	7014.0
32.5	1123.2	1041.3	1180.2	1225.8	1334.0	1242.1	683.3	474.4	653.8	1678.2	6569.6
33.5	1090.1	1024.6	1136.4	1172.5	1264.0	1176.9	660.4	464.8	631.2	1584.6	5918.8
34.5	1049.3	991.4	1086.8	1115.8	1194.9	1108.8	622.6	435.8	585.3	1448.0	5010.9
35.5	1003.4	949.1	1032.1	1054.2	1121.3	1032.9	571.4	389.8	514.8	1257.0	3876.2
36.5	949.9	896.4	967.9	981.5	1034.2	938.9	497.9	316.0	405.0	992.4	2663.2
37.5	894.4	840.3	899.9	903.3	938.9	832.8	411.6	228.0	276.1	700.4	1640.0
38.5	839.8	783.8	830.7	822.5	838.5	720.2	322.2	141.0	154.0	439.0	929.5
39.5	786.1	727.4	761.0	740.0	735.6	606.4	238.9	68.1	58.8	244.5	499.3
40.5	729.9	667.8	686.8	652.3	627.4	490.5	159.4	5.0	16.5	103.5	239.5
41.5	672.4	606.4	610.8	563.8	521.7	383.4	94.1	39.4	63.9	16.9	95.8
42.5	612.7	543.0	533.7	476.6	422.0	288.6	42.1	69.9	93.3	36.1	14.1
43.5	553.2	481.0	460.2	397.1	336.3	213.4	7.4	84.6	104.8	61.2	25.3
44.5	494.5	421.3	392.2	326.9	264.8	155.2	15.7	91.3	108.1	73.1	45.0
45.5	434.4	361.3	326.1	261.1	200.3	103.3	40.1	103.8	118.9	90.5	69.0
46.5	377.0	305.7	267.3	204.8	147.4	61.9	59.6	114.4	128.2	104.5	87.6
47.5	322.9	254.5	214.8	156.0	102.6	27.1	77.8	125.9	138.9	119.1	106.0
48.5	273.1	208.4	168.9	114.1	64.8	2.6	94.7	137.8	150.3	133.8	124.1
49.5	229.1	168.4	130.1	79.4	33.9	26.8	108.8	148.0	160.1	146.2	139.5
50.5	189.4	132.7	95.9	48.9	6.7	48.6	122.8	159.1	171.1	159.7	156.0
51.5	156.3	103.6	68.6	25.0	14.1	64.7	132.3	165.9	177.6	168.0	166.3
52.5	126.5	77.3	44.0	3.4	33.2	80.0	142.3	173.8	185.3	177.6	178.3
53.5	100.0	53.9	22.2	15.9	50.4	94.0	151.7	181.4	192.8	187.1	190.0
54.5	75.2	31.8	1.4	34.7	67.3	108.4	162.2	190.5	202.1	198.6	204.3
55.5	51.5	10.3	19.1	53.4	84.8	123.7	174.2	201.4	213.6	212.5	221.5
56.5	30.6	8.6	36.8	69.6	99.5	136.3	183.6	209.5	221.8	222.5	233.2
57.5	10.5	26.9	54.2	85.5	114.2	149.2	193.5	218.2	230.6	233.1	245.4
58.5	9.5	45.3	71.8	102.0	129.6	163.0	204.6	228.4	241.1	245.5	259.3

Notes: Shows (in thousands) the absolute value of $\widehat{B}_{59.5} - \widehat{M}_{59.5}$ for each possible combination of $a_{59.5,-}$ and $a_{59.5,+}$ for our estimates of the number of taxpayers who changed when they took their first IRA distribution in response to the age thresholds.

Table 14: Robustness check: $\widehat{B}_{59.5} - \widehat{M}_{59.5}$ for reduced-form analysis: gross distributions, first distribution ($bil)

	$a_{59.5,+}$										
$a_{59.5,-}$	59.5	60.5	61.5	62.5	63.5	64.5	65.5	66.5	67.5	68.5	69.5
25.5	7485.5	10.0	8.0	7.8	8061.2	21959.4	6154.0	21.9	23.0	23.5	110388.1
26.5	383.4	7.7	3.2	1.4	677.4	2002.8	17.0	21.2	21.8	16.9	3948.8
27.5	28.6	3.8	12.5	18.6	47.8	70.4	12.5	19.0	18.1	61.8	1813.2
28.5	28.7	6.3	16.1	22.9	40.9	48.1	5.2	1.2	18.9	48.8	89.8
29.5	28.9	13.4	23.1	27.8	35.1	36.8	21.2	20.3	30.7	46.9	73.5
30.5	28.6	20.4	26.3	28.8	32.4	33.1	24.0	23.5	31.3	44.9	68.3
31.5	28.2	23.6	27.2	28.8	31.2	31.7	24.5	24.2	31.0	43.5	65.0
32.5	27.3	24.1	26.8	28.0	29.8	30.1	23.8	23.4	29.7	41.3	60.3
33.5	26.4	23.9	26.1	27.0	28.6	28.8	22.8	22.4	28.2	38.7	54.1
34.5	25.5	23.4	25.2	26.0	27.3	27.4	21.8	21.2	26.4	35.6	45.6
35.5	24.4	22.4	24.0	24.6	25.7	25.6	20.1	19.3	23.5	30.7	33.6
36.5	23.3	21.4	22.7	23.2	24.0	23.6	18.2	17.0	20.1	24.8	21.6
37.5	22.1	20.2	21.4	21.6	22.1	21.4	15.9	14.3	16.1	18.3	11.4
38.5	20.9	19.0	19.9	19.9	20.0	18.8	13.4	11.3	11.8	12.1	4.2
39.5	19.6	17.7	18.3	17.9	17.6	16.0	10.6	8.1	7.7	6.7	0.6
40.5	18.1	16.2	16.5	15.8	15.0	12.9	7.8	5.1	4.1	2.6	3.7
41.5	16.8	14.8	14.7	13.7	12.5	10.3	5.5	3.0	1.9	0.3	4.7
42.5	15.2	13.2	12.8	11.5	10.1	7.7	3.5	1.1	0.0	1.5	5.6
43.5	13.7	11.6	11.0	9.5	7.9	5.6	2.0	0.1	1.1	2.4	5.8
44.5	12.1	10.0	9.2	7.6	6.0	3.9	0.7	1.0	2.0	3.1	6.1
45.5	10.6	8.5	7.6	6.0	4.4	2.5	0.2	1.7	2.6	3.6	6.2
46.5	9.0	7.1	6.0	4.5	3.0	1.2	1.1	2.5	3.3	4.2	6.6
47.5	7.6	5.7	4.6	3.2	1.8	0.2	1.8	3.0	3.8	4.7	6.8
48.5	6.3	4.6	3.5	2.2	0.9	0.6	2.4	3.5	4.2	5.0	7.0
49.5	5.2	3.6	2.6	1.3	0.1	1.1	2.8	3.8	4.4	5.2	7.0
50.5	4.3	2.7	1.7	0.6	0.5	1.7	3.2	4.1	4.7	5.4	7.2
51.5	3.5	2.1	1.1	0.1	0.9	2.0	3.3	4.2	4.8	5.5	7.0
52.5	2.8	1.5	0.6	0.4	1.3	2.3	3.5	4.3	4.9	5.5	7.0
53.5	2.2	1.0	0.2	0.7	1.6	2.5	3.7	4.4	4.9	5.5	6.8
54.5	1.6	0.5	0.3	1.1	1.9	2.8	3.9	4.6	5.1	5.6	6.8
55.5	1.1	0.0	0.7	1.5	2.3	3.1	4.1	4.8	5.3	5.8	6.9
56.5	0.6	0.4	1.1	1.8	2.6	3.3	4.3	4.9	5.3	5.8	6.8
57.5	0.2	0.8	1.4	2.2	2.9	3.6	4.5	5.1	5.5	6.0	6.9
58.5	0.3	1.2	1.8	2.5	3.2	3.9	4.7	5.2	5.7	6.1	6.9

Notes: Shows (in billions) the absolute value of $\widehat{B}_{59.5} - \widehat{M}_{59.5}$ for each possible combination of $a_{59.5,-}$ and $a_{59.5,+}$ for our estimates of the value of the gross distributions taken by taxpayers who changed when they took their first IRA distribution in response to the age thresholds.

Table 15: Robustness check: $\widehat{B}_{59.5} - \widehat{M}_{59.5}$ for reduced-form analysis: number of taxpayers, single year (proportion)

	$a_{59.5,+}$										
$a_{59.5,-}$	59.5	60.5	61.5	62.5	63.5	64.5	65.5	66.5	67.5	68.5	69.5
25.5	3.6234	3.6369	3.5574	3.3918	3.0002	2.4033	2.5100	3.2003	4.2430	6.0148	10.4817
26.5	0.1523	0.3600	0.5712	0.7627	0.9371	1.0700	1.1355	1.1857	1.2524	1.3824	10.4817
27.5	0.8126	0.9828	1.1586	1.3227	1.4772	1.6005	1.6673	1.7219	1.7942	1.9335	48.8059
28.5	1.0516	1.2062	1.3677	1.5207	1.6671	1.7862	1.8531	1.9092	1.9842	2.1297	25.1869
29.5	1.1796	1.3254	1.4788	1.6255	1.7670	1.8835	1.9503	2.0074	2.0843	2.2349	15.7760
30.5	1.2365	1.3760	1.5234	1.6651	1.8026	1.9166	1.9826	2.0395	2.1166	2.2684	9.9137
31.5	1.2662	1.4006	1.5432	1.6809	1.8150	1.9265	1.9914	2.0476	2.1239	2.2740	6.6321
32.5	1.2727	1.4028	1.5411	1.6751	1.8057	1.9146	1.9778	2.0323	2.1060	2.2496	4.3943
33.5	1.2753	1.4017	1.5363	1.6668	1.7943	1.9005	1.9619	2.0145	2.0851	2.2204	3.1488
34.5	1.2681	1.3909	1.5220	1.6491	1.7733	1.8766	1.9356	1.9854	2.0512	2.1733	2.2295
35.5	1.2607	1.3803	1.5079	1.6318	1.7527	1.8530	1.9096	1.9564	2.0169	2.1243	1.6922
36.5	1.2568	1.3734	1.4979	1.6187	1.7365	1.8340	1.8882	1.9321	1.9875	2.0808	1.4344
37.5	1.2398	1.3531	1.4741	1.5913	1.7055	1.7992	1.8500	1.8891	1.9358	2.0057	1.0613
38.5	1.2228	1.3329	1.4504	1.5641	1.6745	1.7644	1.8114	1.8454	1.8829	1.9281	0.8241
39.5	1.2137	1.3209	1.4353	1.5457	1.6525	1.7389	1.7827	1.8125	1.8426	1.8684	0.7739
40.5	1.2001	1.3043	1.4153	1.5221	1.6252	1.7076	1.7476	1.7722	1.7930	1.7957	0.6955
41.5	1.1781	1.2789	1.3861	1.4890	1.5875	1.6652	1.7001	1.7179	1.7265	1.6995	0.5596
42.5	1.1523	1.2496	1.3529	1.4514	1.5451	1.6174	1.6469	1.6569	1.6520	1.5936	0.4340
43.5	1.1232	1.2169	1.3159	1.4099	1.4983	1.5650	1.5883	1.5899	1.5705	1.4803	0.3215
44.5	1.0951	1.1851	1.2799	1.3692	1.4525	1.5134	1.5307	1.5241	1.4911	1.3732	0.2520
45.5	1.0685	1.1548	1.2455	1.3302	1.4082	1.4635	1.4750	1.4606	1.4154	1.2750	0.2167
46.5	1.0400	1.1226	1.2088	1.2888	1.3614	1.4108	1.4163	1.3942	1.3370	1.1766	0.1838
47.5	1.0095	1.0882	1.1700	1.2450	1.3120	1.3555	1.3549	1.3250	1.2563	1.0787	0.1525
48.5	0.9724	1.0469	1.1238	1.1935	1.2544	1.2913	1.2839	1.2453	1.1643	0.9692	0.0976
49.5	0.9327	1.0029	1.0749	1.1390	1.1938	1.2239	1.2097	1.1629	1.0705	0.8613	0.0459
50.5	0.8672	0.9319	0.9976	1.0547	1.1013	1.1223	1.0981	1.0385	0.9281	0.6935	0.1064
51.5	0.7967	0.8559	0.9151	0.9650	1.0032	1.0150	0.9810	0.9092	0.7825	0.5274	0.2472
52.5	0.7117	0.7648	0.8169	0.8588	0.8880	0.8896	0.8447	0.7596	0.6155	0.3396	0.4151
53.5	0.6208	0.6676	0.7125	0.7464	0.7664	0.7581	0.7027	0.6051	0.4457	0.1542	0.5734
54.5	0.5261	0.5666	0.6044	0.6304	0.6416	0.6237	0.5586	0.4500	0.2782	0.0232	0.7172
55.5	0.4244	0.4586	0.4891	0.5071	0.5095	0.4822	0.4078	0.2891	0.1067	0.2012	0.8606
56.5	0.3180	0.3459	0.3693	0.3795	0.3734	0.3372	0.2543	0.1268	0.0638	0.3742	0.9974
57.5	0.2146	0.2365	0.2533	0.2564	0.2426	0.1988	0.1091	0.0246	0.2198	0.5270	1.1095
58.5	0.1059	0.1219	0.1321	0.1283	0.1072	0.0562	0.0397	0.1787	0.3769	0.6794	1.2236

Notes: Shows the absolute value of $\widehat{B}_{59.5} - \widehat{M}_{59.5}$ for each possible combination of $a_{59.5,-}$ and $a_{59.5,+}$ for our estimates of the value of the proportion of taxpayers who changed when they took a distribution from a traditional IRA in response to the age thresholds.

Table 16: Robustness check: $\widehat{B}_{59.5} - \widehat{M}_{59.5}$ for reduced-form analysis: gross distributions, single year ($mil)

	$a_{59.5,+}$										
$a_{59.5,-}$	59.5	60.5	61.5	62.5	63.5	64.5	65.5	66.5	67.5	68.5	69.5
25.5	19.474	28.539	28.933	30.573	77.559	123.478	151.917	252.767	253.322	79.845	2.047
26.5	24.004	24.751	24.675	24.646	25.926	26.008	24.918	24.041	21.156	18.426	2.047
27.5	25.321	25.909	25.874	25.873	26.957	27.044	26.082	25.283	22.542	19.906	145.757
28.5	25.300	25.817	25.784	25.782	26.784	26.863	25.941	25.166	22.485	19.890	49.319
29.5	25.369	25.850	25.820	25.818	26.775	26.851	25.956	25.200	22.570	20.021	37.169
30.5	24.854	25.300	25.260	25.248	26.161	26.222	25.334	24.574	21.960	19.385	6.531
31.5	24.323	24.741	24.691	24.668	25.543	25.588	24.700	23.930	21.321	18.694	5.961
32.5	23.757	24.149	24.088	24.051	24.891	24.917	24.024	23.236	20.619	17.904	12.496
33.5	23.192	23.560	23.487	23.436	24.240	24.245	23.342	22.530	19.895	17.064	15.374
34.5	22.618	22.963	22.877	22.811	23.578	23.559	22.642	21.799	19.134	16.157	16.697
35.5	22.081	22.403	22.304	22.223	22.954	22.909	21.977	21.100	18.401	15.272	16.248
36.5	21.642	21.945	21.835	21.741	22.438	22.371	21.425	20.520	17.800	14.556	13.695
37.5	21.090	21.370	21.245	21.132	21.789	21.691	20.721	19.769	16.998	13.560	13.516
38.5	20.717	20.979	20.844	20.718	21.342	21.222	20.241	19.262	16.481	12.970	10.446
39.5	20.308	20.550	20.404	20.263	20.851	20.703	19.707	18.694	15.894	12.291	8.505
40.5	19.935	20.160	20.002	19.848	20.400	20.227	19.218	18.175	15.370	11.717	6.452
41.5	19.471	19.674	19.501	19.328	19.837	19.631	18.597	17.509	14.672	10.911	5.714
42.5	18.871	19.047	18.853	18.654	19.110	18.859	17.784	16.625	13.717	9.759	6.229
43.5	18.454	18.609	18.401	18.185	18.600	18.319	17.227	16.035	13.130	9.173	4.886
44.5	17.803	17.928	17.695	17.449	17.806	17.472	16.335	15.064	12.088	7.963	5.494
45.5	17.430	17.536	17.293	17.032	17.351	16.993	15.849	14.565	11.635	7.631	3.882
46.5	16.914	16.995	16.733	16.451	16.720	16.324	15.157	13.833	10.909	6.932	3.422
47.5	16.267	16.318	16.032	15.719	15.930	15.484	14.279	12.893	9.944	5.933	3.706
48.5	15.625	15.646	15.336	14.995	15.148	14.656	13.419	11.983	9.030	5.040	3.813
49.5	15.110	15.108	14.782	14.422	14.529	14.007	12.761	11.312	8.422	4.586	3.208
50.5	13.921	13.866	13.489	13.067	13.076	12.456	11.112	9.519	6.479	2.386	5.651
51.5	12.900	12.800	12.383	11.913	11.842	11.150	9.742	8.062	4.970	0.825	6.897
52.5	11.524	11.364	10.890	10.353	10.178	9.383	7.875	6.056	2.838	1.486	9.229
53.5	10.218	10.005	9.481	8.885	8.618	7.738	6.154	4.238	0.960	3.410	10.885
54.5	8.715	8.442	7.860	7.197	6.828	5.854	4.183	2.156	1.196	5.631	12.914
55.5	6.977	6.636	5.987	5.247	4.768	3.687	1.916	0.234	3.678	8.199	15.357
56.5	5.270	4.864	4.156	3.347	2.768	1.597	0.251	2.490	5.972	10.488	17.377
57.5	3.497	3.028	2.260	1.385	0.711	0.544	2.460	4.772	8.268	12.745	19.335
58.5	1.560	1.024	0.192	0.754	1.526	2.867	4.853	7.237	10.750	15.189	21.520

Notes: Shows (in millions) the absolute value of $\widehat{B}_{59.5} - \widehat{M}_{59.5}$ for each possible combination of $a_{59.5,-}$ and $a_{59.5,+}$ for our estimates of change of the value of the average distribution from traditional IRAs in response to the age thresholds.

E Intuition from a stylized model

(For online publication)

In this section, we provide intuition for why we would expect taxpayers to bunch at the age thresholds in order to avoid the early withdrawal penalty. We consider a simple, highly-stylized model that allows us to graphically demonstrate our anticipated bunching. We also discuss how we can use bunching in this setting to identify the elasticity of intertemporal substitution.

E.1 A simple model with no penalties

Consider a taxpayer with assets A in an account earning a rate of return $R = 1 + r$. These assets are the only source of wealth available to the taxpayer, and the taxpayer must withdraw all of the assets at once. Once the assets are withdrawn, they stop earning the rate of return and the taxpayer's lifetime consumption is then equal to the value of the account. The taxpayer chooses when to withdraw the assets and their subsequent consumption stream in order to maximize her lifetime utility from consumption, as shown in Equation 24:

$$\max_{\{c_t\},t^*} \sum_{t=0}^{T} \beta^t u(c_t) \tag{24}$$

$$\text{subject to}$$

$$\sum_{t=t^*}^{T} c_t = R^{t^*} A$$

The benefit of withdrawing in an earlier time period is the taxpayer receives utility from consumption earlier (the discount factor does not have as much of an impact on lifetime utility). The benefit of waiting to withdraw is the account grows in value and total lifetime consumption is higher. The taxpayer will pick t^* to equalize the marginal benefit of these two options, given their initial level of assets A and the rate of return R.

E.2 Anticipated bunching due to an early withdrawal penalty

Now introduce a penalty ρ on withdrawing before $t = t_p$. This reduces the value of lifetime consumption for withdrawing before $t = t_p$ by ρ times the values of the assets at the point of withdrawal. The taxpayer's new problem is shown in Equation 25:

$$\max_{t^*} \sum_{t=0}^{T} \beta^t u(c_t) \tag{25}$$

subject to

$$\sum_{t=t^*}^{T} c_t = R^{t^*} A - \rho R^{t^*} A \mathbb{I}(t^* < t_p)$$

Some taxpayers who otherwise would have chosen $t^* < t_p$ will wait to withdraw until after $t = t_p$.

Figure 17 shows the choice of a taxpayer deciding when to withdraw assets from an account when faced with an early withdrawal penalty and the subsequent bunching in the distribution of chosen ages for withdrawal. Figure 17a illustrates how the taxpayer's choice of when to withdraw the assets from the account translates directly into total lifetime consumption. The solid black line represents the budget constraint when there is a penalty for withdrawing the assets before age t_p. The slope of the budget constraint is equal to $tR^{t-1}A$ above t_p and $(1-\rho)tR^{t-1}A$ below. The discrete change in the budget constraint at $t = t_p$ is due to the lack of early withdrawal penalty above that age. The solid red line represents the indifference curve for the marginal buncher. The marginal buncher is indifferent to withdrawing at ages t_I and t_p. The dashed red line represents the indifference curve for the marginal buncher in the absence of the early withdrawal penalty. In this scenario, the taxpayer would have chosen t^*.

Figure 17b demonstrates how the early withdrawal penalty might impact the distribution of withdrawal ages chosen. We've drawn a potential, arbitrary distribution as a dotted red line. The solid red line illustrates what happens when we introduce the early withdrawal penalty. Taxpayers for whom the optimal age of withdrawal would have been between t_I and t_p now prefer to wait to withdraw until t_p. This results in bunching in the distribution of withdrawal ages at t_p.

E.3 Identifying the elasticity of intertemporal substitution

The elasticity of intertemporal substitution (EIS) measures the relative response of consumption in period 1 versus consumption in period 2 when the rate of return changes. When the real rate of return increases, two things happen. First, the taxpayer's lifetime budget increases. This suggests an increase in consumption in all periods (an income effect). Second, consumption in period 1 becomes more expensive, because it means the

Figure 17: Bunching due to early withdrawal penalty in a highly stylized model

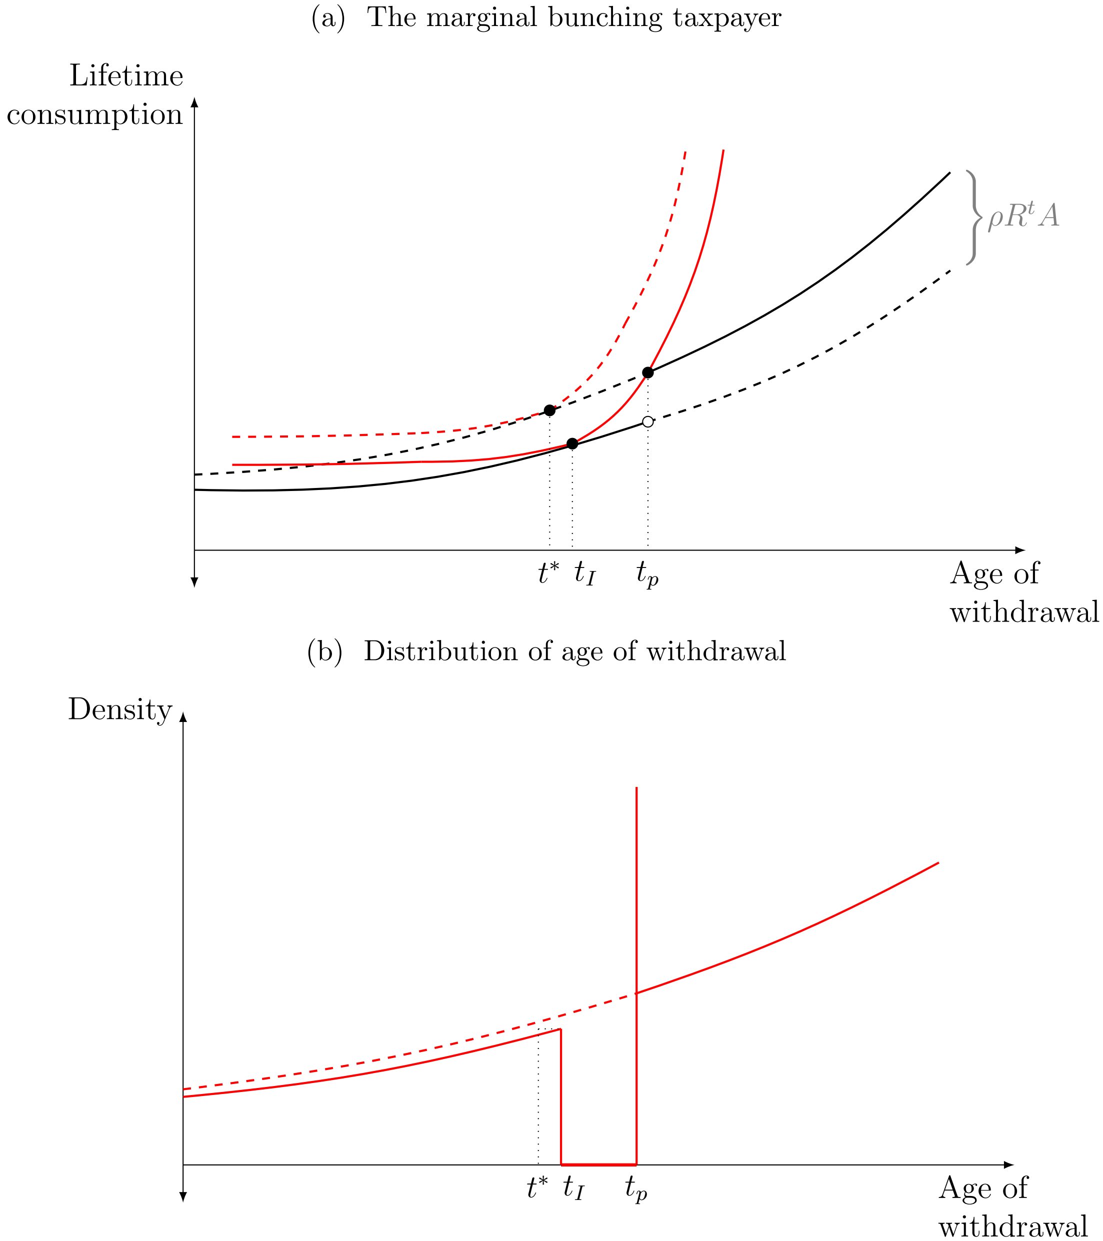

Notes: The distribution of "age of withdrawal" is arbitrary and used for illustrative purposes only.

taxpayer saves less (and therefore benefits less from the increased interest rate). This suggests a decrease of consumption in period 1 in favor of consumption in period 2 (a substitution effect). The net effect of these two forces is represented by the EIS.

The value of the EIS directly impacts the curvature of the indifference curve. Greater values of the EIS equal less curved indifference curves. The more elastic the taxpayer's preferences are between periods 1 and periods 2, the less period 2 consumption the taxpayer needs to be compensated with to maintain their utility level if their period 1 consumption decreases. We can take advantage of the relationship between the EIS and the curvature of the indifference curve to identify the EIS in our setting. Figure 18 shows the same setting as in Figure 17, but where the indifference curves are drawn using different values of the EIS. The taxpayers represented in the left-hand panel have more strongly curved indifference curves (i.e., a lower EIS) than the taxpayers in the right-hand panel. This causes the space between t_I and t_p to be narrower, resulting in fewer taxpayers bunching at t_p. The amount of bunching at t_p, then, is directly related to the value of the EIS.

Figure 18: Identifying the EIS using a bunching response

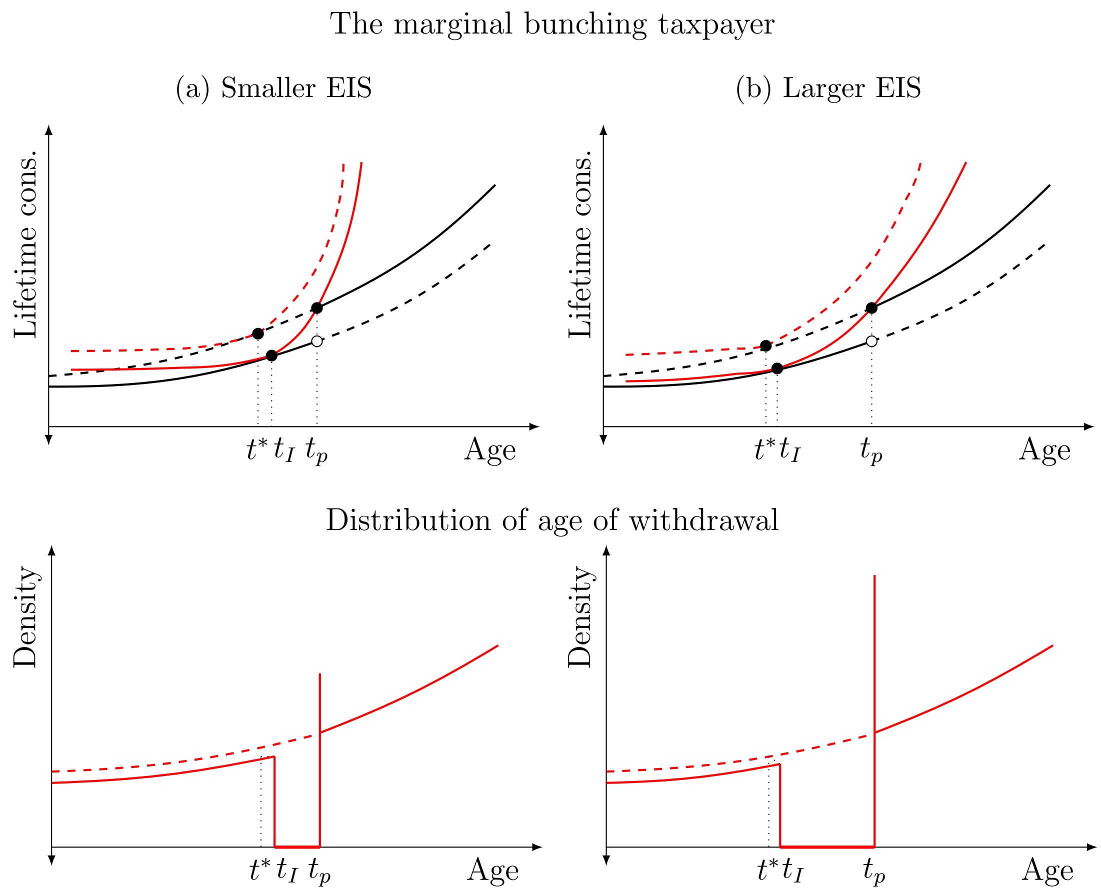

Notes: The distribution of "age of withdrawal" is arbitrary and used for illustrative purposes only.

The previous argument assumed that we knew the value of the discount factor. In the model presented above, present versus future consumption decisions are dictated by two parameters: the EIS and the discount factor. If we only had one notch, we would not separately be able to identify either parameters. However, there are two age notches in our setting: the notch used in the simple model above (equivalent to the notch created by the early withdrawal penalty), and the notch generated by the onset of required minimum distributions. Because we have two bunching conditions, we are able to separately identify the two relevant parameters.

F Parameters used in the structural model

(For online publication)

Table 17 defines the parameters used in the life-cycle model described in Section 4.

Table 17: Structural model parameters

Preference parameters	
β	Discount factor
σ	Elasticity of intertemporal substitution
A	Bequest weight
α	Bequest elasticity
State variables	
Ω	Set of all state variables
S_t	Stock of assets in standard savings account
A_t	Stock of assets in tax-benefited savings account
θ	Labor income shock
z_t	Total income subject to income tax ($= y_t + P \cdot \mathbb{I}[t \geq t_P] + q(S_t \cdot r)$)
t	period (equivalent to age)
P	Value of annual pension
t_P	Age at which individual claims pension
Choice variables	
S_{t+1}	Next period stock of assets in standard savings account
A_{t+1}	Next period stock of assets in tax-benefited savings account
c_t	Consumption
Labor income	
y_t	Labor income
$\{\rho_i\}_{i=0}^3$	Deterministic part of labor income
ε_{it}	Stochastic part of labor income
η	Autocorrelation of stochastic part of labor income
σ_ζ^2	Variance of stochastic term
Tax-benefited savings account parameters	
$\overline{a}_t$	Contribution limit for tax-benefited savings account
t_e	Age before which early withdrawal penalty applies
p_e	Rate of early withdrawal penalty
t_{rmd}	Age after which excess accumulation penalty applies
p_{rmd}	Rate of excess accumulation penalty
$a_{t,rmd}$	Value of RMD (function of age and A_t)
Taxes	
τ_y	Income tax function
τ_e	Amount of early withdrawal penalty
τ_{rmd}	Amount of excess accumulation penalty
Additional parameters	
r_A	Rate of return on regular savings
r_S	Rate of return on tax-benefited savings
π_t	The probability of living to period t conditional on surviving to period $t-1$
W_{t+1}	Bequeathed wealth if die at end of period t ($= (1+r_S)S_{t+1} + (1+r_A)A_{t+1}$)

G Numerical solution

(For online publication)

We numerically solve the dynamic optimization problem presented in Section 4 using backwards iteration. Two of our state variables are already discrete: age, and age at which the individual claims Social Security. We solve the model for three values of claiming Social Security age: 62, 67, and 70. These ages respectively correspond to the earliest age an individual is eligible to claim Social Security, the full retirement age, and the age after which delayed retirement credits stop permanently increasing the monthly (and therefore annual) Social Security benefit for individuals born in 1940 or later.

We have three continuous state variables that we need to discretize: labor income shocks and the stock of both regular savings and tax-benefited savings. Labor income shocks are placed on a 5-point grid using the Rouwenhorst method (Rouwenhorst 1995).

Discretizing the stocks of tax-benefited saving is finicky in our setting. For each period in our model, we need at least one non-zero contribution option, and a "no contribution" option. Ensuring that both of these options were available at each period required that we make some non-standard choices. We space the grid at the oldest age individuals in our model can make a contribution (70) by the contribution limit at age 69. We then create age-specific grids working backwards, where grid point $A_{j,t} = A_{j,t+1}/(1 + r_A)$. Because savings accounts will be depleted as the simulated individuals near the final period, we also create age-specific grids from 70 forward such that $A_{j,t} = A_{j,t-1}/(1 + r_A)$.

This methodology ensures that, for any grid point $A_{j,t}$ with $t < 70$, the model can choose not to contribute by choosing point $A_{j,t+1}$.[44] This procedure also ensures that the model always has at least one gridpoint to choose that is within the contribution limit, because the grid points are spaced within the contribution limit at the oldest age one can make a contribution and become closer together as age decreases. We set the standard savings grid to be equal to the grid used for tax-benefited savings so that, when solving the model, it is never beneficial to take a distribution from the tax-benefited account in order to save in the regular savings account. In other words, the next-period options for both accounts are the same.

In order to not severely truncate the maximum level of assets that can be saved in these accounts, we use 90-point grids for both regular and tax-benefited savings. This is considerably larger than is often used in these models (e.g., Choukhmane (2021) uses 20 grid points for liquid assets and 16 grid points for retirement wealth). Solving the model for ages 40-85 with two 90-point asset grids, a 5-point labor income shock grid, and 3 possible

[44] This is because choosing not to contribute is equivalent to picking a next-period stock equal to the current stock multiplied by the return.

ages for claiming Social Security means that we are finding the solution for 5,589,000 combinations of the state variables, each of which faces 8,100 possible combinations of the choice variables.[45] Even with these large grids, the top 3.7% of IRA balances in our simulated sample are constrained by our highest grid value at age 40, and the top 1.6% are constrained at age 41.

We minimize the score ($Z(\Phi, n)$, see Equation 22 using Nelder-Mead optimization. Because Nelder-Mead can get stuck at local minima, we pick our starting point for the SMM algorithm as follows: we generate 4,000 points defined by a Sobol sequence over a wide range of possible parameter values: between 0.1 and 2.0 for the EIS (σ) and bequest elasticity (α), between 0.85 and 1.0 for the discount factor (β), and between 0.25 and 5.0 for the bequest weight (A). We then calculate the score for all 4,000 preference parameter combinations. We use the parameter combination that minimizes the score as our initial guess for the SMM procedure.

Estimating the model is computationally intensive. The code is written in Python. All functions are JIT-compiled in "nopython" mode using Numba v0.49.1 with parallel = True, cache = True, and all object types defined.

[45] We solve for optimal next-period stocks in both assets using a grid search, which implies post-tax consumption. Not every state space has this many options because of the contribution limit to the tax-benefited savings account and the no borrowing constraint.

H Counterfactual policy analysis

(For online publication)

H.1 Changing the excess accumulation penalty

Table 18 shows the results of changing the excess accumulation penalty. We compare the values of four outcomes to the base policy with an age threshold = $70\frac{1}{2}$ and penalty rate = 50%. For all outcomes considered, a positive value indicates that our simulated individuals are, on average, better off in the counterfactual world, while a negative value indicates that our simulated individuals are, on average, better off in the base policy.

Table 18a shows a measure of equivalent variation: the average amount of income that our simulated individuals would need to receive (give up) in the base policy to reach the present discounted value of lifetime utility for each counterfactual policy. We show the value in levels, and as a percentage of average income at age $40\frac{1}{2}$ (i.e., the first period in our model). Table 18b shows the change in the present discounted value of lifetime total tax remittances relative to the base policy, where total tax remittances include income taxes plus any tax penalties paid. Table 18c shows the change in the present discounted value of of lifetime income tax remittances relative to the base policy. Table 18d shows the change in the bequeathed IRA balance relative to the base policy. For each panel, we show the amount in levels and as a percentage of the value in the base policy.

Table 18: Results of changing the excess accumulation penalty

(a) Change in income at age $40\frac{1}{2}$ under base policy to reach counterfactual PDV lifetime utility

Age	40% penalty		50% penalty		60% penalty		70% penalty	
threshold	Levels ($)	Percent	Levels ($)	Percent	Levels ($)	Percent	Levels ($)	Percent
$68\frac{1}{2}$	-761	-0.0037%	-811	-0.0040%	-851	-0.0042%	-878	-0.0043%
$69\frac{1}{2}$	-335	-0.0016%	-362	-0.0018%	-395	-0.0019%	-411	-0.0020%
$70\frac{1}{2}$	22	0.0001%	0	0.0000%	-38	-0.0002%	-53	-0.0003%
$71\frac{1}{2}$	204	0.0010%	186	0.0009%	168	0.0008%	151	0.0007%
$72\frac{1}{2}$	344	0.0017%	339	0.0017%	325	0.0016%	309	0.0015%
$73\frac{1}{2}$	459	0.0023%	454	0.0022%	451	0.0022%	441	0.0022%
$74\frac{1}{2}$	533	0.0026%	531	0.0026%	530	0.0026%	525	0.0026%
$75\frac{1}{2}$	565	0.0028%	570	0.0028%	570	0.0028%	571	0.0028%
$76\frac{1}{2}$	579	0.0029%	588	0.0029%	595	0.0029%	597	0.0029%
$77\frac{1}{2}$	568	0.0028%	586	0.0029%	597	0.0029%	604	0.0030%
$78\frac{1}{2}$	550	0.0027%	569	0.0028%	579	0.0028%	584	0.0029%

Notes: Each iteration includes 10,000 unique simulated individuals. This table compares the base policy for the excess accumulation policy (age threshold = $70\frac{1}{2}$, penalty rate = 50%) and a range of counterfactual policies. Table 18a shows the average amount of income that our simulated individuals would need to receive (give up) in the base policy to reach the present discounted value of lifetime utility for each counterfactual policy. Table 18b shows the change in the present discounted value of lifetime total tax remittances. Table 18c shows the change in the present discounted value of lifetime income tax remittances. Table 18d shows the change in the bequeathed IRA balance. For each panel, we show the amount in levels and as a percentage of the value in the base policy.

Table 18: Results of changing the excess accumulation penalty, continued

(b) Change in PDV lifetime total tax remittances relative to base policy

Age threshold	40% penalty Levels ($)	40% penalty Percent	50% penalty Levels ($)	50% penalty Percent	60% penalty Levels ($)	60% penalty Percent	70% penalty Levels ($)	70% penalty Percent
$68\frac{1}{2}$	-57	-0.0178%	8	0.0026%	53	0.0165%	77	0.0243%
$69\frac{1}{2}$	-27	-0.0086%	59	0.0187%	122	0.0385%	140	0.0442%
$70\frac{1}{2}$	-94	-0.0296%	0	0.0000%	64	0.0201%	88	0.0276%
$71\frac{1}{2}$	-124	-0.0391%	-38	-0.0120%	25	0.0079%	67	0.0211%
$72\frac{1}{2}$	-167	-0.0526%	-81	-0.0254%	-16	-0.0052%	20	0.0062%
$73\frac{1}{2}$	-225	-0.0707%	-143	-0.0449%	-75	-0.0237%	-32	-0.0102%
$74\frac{1}{2}$	-288	-0.0905%	-206	-0.0649%	-140	-0.0442%	-103	-0.0323%
$75\frac{1}{2}$	-373	-0.1175%	-295	-0.0928%	-236	-0.0743%	-195	-0.0614%
$76\frac{1}{2}$	-478	-0.1504%	-404	-0.1273%	-346	-0.1089%	-311	-0.0980%
$77\frac{1}{2}$	-601	-0.1890%	-534	-0.1682%	-480	-0.1509%	-445	-0.1402%
$78\frac{1}{2}$	-731	-0.2299%	-672	-0.2115%	-624	-0.1964%	-596	-0.1875%

(c) Change in PDV lifetime income tax remittances relative to base policy

Age threshold	40% penalty Levels ($)	40% penalty Percent	50% penalty Levels ($)	50% penalty Percent	60% penalty Levels ($)	60% penalty Percent	70% penalty Levels ($)	70% penalty Percent
$68\frac{1}{2}$	-205	-0.0648%	-73	-0.0232%	50	0.0157%	156	0.0495%
$69\frac{1}{2}$	-109	-0.0344%	13	0.0040%	149	0.0473%	250	0.0792%
$70\frac{1}{2}$	-123	-0.0388%	0	0.0000%	137	0.0434%	240	0.0761%
$71\frac{1}{2}$	-98	-0.0310%	33	0.0104%	161	0.0509%	240	0.0758%
$72\frac{1}{2}$	-95	-0.0300%	27	0.0086%	142	0.0451%	234	0.0740%
$73\frac{1}{2}$	-97	-0.0307%	25	0.0078%	128	0.0405%	205	0.0649%
$74\frac{1}{2}$	-125	-0.0395%	-7	-0.0021%	95	0.0301%	167	0.0528%
$75\frac{1}{2}$	-159	-0.0504%	-56	-0.0176%	39	0.0124%	98	0.0311%
$76\frac{1}{2}$	-212	-0.0670%	-125	-0.0396%	-44	-0.0138%	16	0.0050%
$77\frac{1}{2}$	-285	-0.0901%	-208	-0.0659%	-133	-0.0421%	-84	-0.0265%
$78\frac{1}{2}$	-377	-0.1192%	-309	-0.0978%	-239	-0.0757%	-198	-0.0627%

(d) Bequeathed IRA balance relative to base policy

Age threshold	40% penalty Levels ($)	40% penalty Percent	50% penalty Levels ($)	50% penalty Percent	60% penalty Levels ($)	60% penalty Percent	70% penalty Levels ($)	70% penalty Percent
$68\frac{1}{2}$	-5,580	-8.6%	-7,286	-11.3%	-8,910	-13.8%	-10,176	-15.8%
$69\frac{1}{2}$	-1,937	-3.0%	-3,527	-5.5%	-5,159	-8.0%	-6,462	-10.0%
$70\frac{1}{2}$	1,441	2.2%	0	0.0%	-1,539	-2.4%	-2,863	-4.4%
$71\frac{1}{2}$	3,554	5.5%	1,992	3.1%	516	0.8%	-352	-0.5%
$72\frac{1}{2}$	5,724	8.9%	4,283	6.6%	2,946	4.6%	1,976	3.1%
$73\frac{1}{2}$	7,504	11.6%	6,198	9.6%	5,045	7.8%	4,297	6.7%
$74\frac{1}{2}$	9,360	14.5%	8,130	12.6%	7,041	10.9%	6,330	9.8%
$75\frac{1}{2}$	11,053	17.1%	10,004	15.5%	9,022	14.0%	8,397	13.0%
$76\frac{1}{2}$	12,642	19.6%	11,782	18.3%	10,926	16.9%	10,309	16.0%
$77\frac{1}{2}$	14,246	22.1%	13,454	20.9%	12,657	19.6%	12,142	18.8%
$78\frac{1}{2}$	15,788	24.5%	15,106	23.4%	14,417	22.3%	14,033	21.7%

Notes: Each iteration includes 10,000 unique simulated individuals. This table compares the base policy for the excess accumulation policy (age threshold = $70\frac{1}{2}$, penalty rate = 50%) and a range of counterfactual policies. Table 18a shows the average amount of income that our simulated individuals would need to receive (give up) in the base policy to reach the present discounted value of lifetime utility for each counterfactual policy. Table 18b shows the change in the present discounted value of lifetime total tax remittances. Table 18c shows the change in the present discounted value of lifetime income tax remittances. Table 18d shows the change in the bequeathed IRA balance. For each panel, we show the amount in levels and as a percentage of the value in the base policy.

H.2 Changing the early withdrawal penalty

Table 19 shows the results of changing the early withdrawal penalty. We compare the values of four outcomes to the base policy with an age threshold = $59\frac{1}{2}$ and penalty rate = 10%. For all outcomes considered, a positive value indicates that our simulated individuals are, on average, better off in the counterfactual world, while a negative value indicates that our simulated individuals are, on average, better off in the base policy.

Table 19a shows a measure of equivalent variation: the average amount of income that our simulated individuals would need to receive (give up) in the base policy to reach the present discounted value of lifetime utility for each counterfactual policy. We show the value in levels, and as a percentage of average income at age $40\frac{1}{2}$ (i.e., the first period in our model). Table 19b shows the change in the present discounted value of lifetime total tax remittances relative to the base policy, where total tax remittances include income taxes plus any tax penalties paid. Table 19c shows the change in the present discounted value of of lifetime income tax remittances relative to the base policy. Table 19d shows the change in IRA balance at age $65\frac{1}{2}$ relative to the base policy. For each panel, we show the amount in levels and as a percentage of the value in the base policy.

Table 19: Results of changing the early withdrawal penalty

(a) Change in income at age $40\frac{1}{2}$ under base policy to reach counterfactual PDV lifetime utility

Age	5% penalty		10% penalty		20% penalty		30% penalty	
threshold	Levels ($)	Percent	Levels ($)	Percent	Levels ($)	Percent	Levels ($)	Percent
$55\frac{1}{2}$	3,427	0.0169%	1,408	0.0069%	-1,600	-0.0079%	-3,806	-0.0187%
$56\frac{1}{2}$	3,246	0.0160%	1,092	0.0054%	-2,140	-0.0105%	-4,472	-0.0220%
$57\frac{1}{2}$	3,076	0.0151%	749	0.0037%	-2,685	-0.0132%	-5,205	-0.0256%
$58\frac{1}{2}$	2,884	0.0142%	390	0.0019%	-3,302	-0.0162%	-5,964	-0.0293%
$59\frac{1}{2}$	2,672	0.0131%	0	0.0000%	-3,892	-0.0191%	-6,753	-0.0332%
$60\frac{1}{2}$	2,462	0.0121%	-365	-0.0018%	-4,521	-0.0222%	-7,547	-0.0371%
$61\frac{1}{2}$	2,237	0.0110%	-724	-0.0036%	-5,205	-0.0256%	-8,404	-0.0413%
$62\frac{1}{2}$	2,004	0.0099%	-1,164	-0.0057%	-5,921	-0.0291%	-9,353	-0.0460%
$63\frac{1}{2}$	1,862	0.0092%	-1,488	-0.0073%	-6,445	-0.0317%	-10,034	-0.0494%
$64\frac{1}{2}$	1,684	0.0083%	-1,817	-0.0089%	-7,026	-0.0346%	-10,716	-0.0527%
$65\frac{1}{2}$	1,486	0.0073%	-2,148	-0.0106%	-7,665	-0.0377%	-11,519	-0.0567%

Notes: Each iteration includes 10,000 unique simulated individuals. This table compares the base policy for the early withdrawal policy (age threshold = $59\frac{1}{2}$, penalty rate = 10%) and a range of counterfactual policies. Table 19a shows the average amount of income that our simulated individuals would need to receive (give up) in the base policy to reach the present discounted value of lifetime utility for each counterfactual policy. Table 19b shows the change in the present discounted value of lifetime total tax remittances. Table 19c shows the change in the present discounted value of lifetime income tax remittances. Table 19d shows the change in IRA balances at age $65\frac{1}{2}$. For each panel, we show the amount in levels and as a percentage of the value in the base policy.

Table 19: Results of changing the early withdrawal penalty, continued

(b) Change in PDV lifetime total tax remittances relative to base policy

Age	5% penalty		10% penalty		20% penalty		30% penalty	
threshold	Levels ($)	Percent	Levels ($)	Percent	Levels ($)	Percent	Levels ($)	Percent
$55\frac{1}{2}$	-385	-0.1210%	-167	-0.0525%	20	0.0062%	144	0.0453%
$56\frac{1}{2}$	-368	-0.1159%	-160	-0.0505%	-7	-0.0021%	94	0.0297%
$57\frac{1}{2}$	-341	-0.1072%	-132	-0.0415%	-20	-0.0062%	33	0.0105%
$58\frac{1}{2}$	-301	-0.0948%	-73	-0.0231%	4	0.0014%	-4	-0.0012%
$59\frac{1}{2}$	-250	-0.0787%	0	0.0000%	78	0.0245%	-3	-0.0009%
$60\frac{1}{2}$	-179	-0.0563%	100	0.0316%	199	0.0625%	70	0.0221%
$61\frac{1}{2}$	-84	-0.0264%	244	0.0769%	369	0.1162%	225	0.0709%
$62\frac{1}{2}$	41	0.0130%	432	0.1360%	620	0.1951%	449	0.1413%
$63\frac{1}{2}$	144	0.0453%	597	0.1880%	844	0.2657%	669	0.2104%
$64\frac{1}{2}$	268	0.0845%	801	0.2520%	1,138	0.3580%	937	0.2948%
$65\frac{1}{2}$	412	0.1295%	1,057	0.3326%	1,512	0.4756%	1,300	0.4092%

(c) Change in PDV lifetime income tax remittances relative to base policy

Age	5% penalty		10% penalty		20% penalty		30% penalty	
threshold	Levels ($)	Percent	Levels ($)	Percent	Levels ($)	Percent	Levels ($)	Percent
$55\frac{1}{2}$	521	0.1649%	379	0.1199%	129	0.0407%	28	0.0088%
$56\frac{1}{2}$	470	0.1489%	274	0.0867%	-51	-0.0161%	-188	-0.0594%
$57\frac{1}{2}$	424	0.1342%	178	0.0564%	-236	-0.0747%	-429	-0.1358%
$58\frac{1}{2}$	379	0.1199%	89	0.0280%	-422	-0.1334%	-684	-0.2167%
$59\frac{1}{2}$	332	0.1052%	0	0.0000%	-586	-0.1856%	-933	-0.2952%
$60\frac{1}{2}$	292	0.0924%	-90	-0.0284%	-746	-0.2363%	-1,155	-0.3655%
$61\frac{1}{2}$	254	0.0804%	-175	-0.0554%	-903	-0.2859%	-1,363	-0.4316%
$62\frac{1}{2}$	216	0.0683%	-262	-0.0830%	-1,064	-0.3367%	-1,579	-0.4997%
$63\frac{1}{2}$	182	0.0576%	-331	-0.1048%	-1,192	-0.3772%	-1,738	-0.5501%
$64\frac{1}{2}$	150	0.0476%	-408	-0.1293%	-1,325	-0.4196%	-1,939	-0.6139%
$65\frac{1}{2}$	108	0.0343%	-486	-0.1539%	-1,464	-0.4635%	-2,153	-0.6814%

(d) IRA balance at age $65\frac{1}{2}$ relative to base policy

Age	5% penalty		10% penalty		20% penalty		30% penalty	
threshold	Levels ($)	Percent	Levels ($)	Percent	Levels ($)	Percent	Levels ($)	Percent
$55\frac{1}{2}$	-186	-0.1912%	-40	-0.0413%	31	0.0322%	183	0.1875%
$56\frac{1}{2}$	-209	-0.2146%	-106	-0.1084%	-83	-0.0847%	68	0.0695%
$57\frac{1}{2}$	-221	-0.2271%	-121	-0.1247%	-177	-0.1816%	-41	-0.0426%
$58\frac{1}{2}$	-202	-0.2072%	-106	-0.1084%	-192	-0.1969%	-122	-0.1257%
$59\frac{1}{2}$	-138	-0.1421%	0	0.0000%	-121	-0.1246%	-104	-0.1071%
$60\frac{1}{2}$	-39	-0.0396%	143	0.1463%	106	0.1090%	161	0.1653%
$61\frac{1}{2}$	135	0.1386%	384	0.3942%	520	0.5335%	655	0.6722%
$62\frac{1}{2}$	322	0.3304%	748	0.7674%	1,148	1.1781%	1,498	1.5375%
$63\frac{1}{2}$	484	0.4969%	1,126	1.1561%	1,804	1.8519%	2,476	2.5419%
$64\frac{1}{2}$	712	0.7311%	1,572	1.6139%	2,666	2.7370%	3,646	3.7424%
$65\frac{1}{2}$	1,096	1.1249%	2,279	2.3391%	3,931	4.0346%	5,309	5.4494%

Notes: Each iteration includes 10,000 unique simulated individuals. This table compares the base policy for the early withdrawal policy (age threshold = $59\frac{1}{2}$, penalty rate = 10%) and a range of counterfactual policies. Table 19a shows the average amount of income that our simulated individuals would need to receive (give up) in the base policy to reach the present discounted value of lifetime utility for each counterfactual policy. Table 19b shows the change in the present discounted value of lifetime total tax remittances. Table 19c shows the change in the present discounted value of lifetime income tax remittances. Table 19d shows the change in IRA balances at age $65\frac{1}{2}$. For each panel, we show the amount in levels and as a percentage of the value in the base policy.