Understanding the influence of marginal income tax rates on retirement investment habits *

Daniel Cohen†

March 5, 2023

Abstract

I find via difference-in-differences analysis of IRS filings between 1999 and 2018 that an increase in marginal income tax rate is associated with increases in retirement plan contributions and participation among employees, two key metrics of retirement preparedness. This discovery expands upon the prevailing literature, which establishes the sensitivity of taxable income to taxation and affirms the ability of federal and state governments to enact legislation that affects private investment habits. My findings suggest that increasing income tax rates may help address the United States retirement unpreparedness crisis by further incentivizing retirement savings.

*This research is funded in part by the UC Berkeley Summer Undergraduate Research Fellowships program and its private donors. I thank Dr. Danny Yagan and Dr. Gabriel Zucman, my primary and secondary thesis advisors, for their advice and suggestions throughout the research process. I also thank Dr. Hunt Allcott for his valuable input and guidance.

†Department of Economics at the University of California, Berkeley, daniel.m.cohen@berkeley.edu.

Contents

1 Introduction

Americans habitually under-save for retirement, with nearly half of all adults aged 55+ lacking any form of retirement savings (GAO 2015). There is a significant body of research dedicated to understanding why this is the case: economists have studied the role of peer effects (Beshears et al. 2015), nudges (Beshears et al. 2021), and present bias (Beshears et al. 2022) as possible explanations of (or solutions to) this phenomenon. However, these auxiliary considerations all exist in the shadow of a progressive income tax, which provides the primary incentive for investment in tax-deferred retirement accounts. Hence it is important to answer the following question: What is the elasticity of retirement investing with respect to marginal income tax rates? Answering this question will help contextualize the impact of proposed solutions to retirement under-saving, and additionally reveal whether or not income tax increases themselves can incentivize retirement savings.

Intuitively, there are two opposing forces at play when income tax rates are increased. The first is a substitution effect, whereby increasing the marginal income tax rate causes the relative price of retirement investment to fall, thus incentivizing additional savings. However, there is also an opposing income effect: if a tax increase is sufficiently large, it may result in income reductions that render retirement investing impossible or suboptimal. This is particularly true if consumers are myopic or heavily biased. As a result, it is necessary to conduct an empirical analysis to understand which of these effects dominates, and hence determine the elasticity of retirement investing with respect to income tax rates.

This question is highly relevant from a policy and welfare perspective. Because it is anticipated that Social Security fund depletion will cause significant cuts to benefits by 2037 (Goss 2010), private pensions must pick up an increasing amount of slack if full retirement income is to be ensured for US private employees. However, private pensions themselves are often under-supported, providing further barriers to saving. Soto (2008) explains that volatility in private pension costs has approached unsustainable levels for even large US employers, who are now dialing down their retirement support. Andrews (1989) elaborates upon the chronic lack of small-business private pensions in the US. Considering this deterioration of, and under-contribution to, existing private pension plans, it is imperative to understand how existing policy levers (such as income taxation) can remedy this issue and encourage increased retirement savings.

In this paper, I study the influence of marginal income tax rates on private pension investments. I use IRS Form 5500 filings to estimate difference-in-differences regressions for various retirement investment outcomes, where treatment is akin to a large, isolated increase in marginal state income tax rate. My analysis considers all US "small businesses" (per IRS definitions) between 1999 and 2018. My approach is unique because the use of Form 5500 data allows me to capture a large portion of US private employees and observe both firm and consumer investment decisions.

In these regressions of retirement investment outcomes against treatment, controlling for state-specific macroeconomic and pension plan characteristics (GDP, income, education, population, employment, and plan age), I find increases in the marginal income tax rate are associated with significant increases in private pension plan participation, contributions, and net assets. However, my robustness checks are inconclusive, and strategically tailoring my regression specification reduces the significance of my findings. My results elaborate upon the prevailing literature, which has found strong relationships between taxation and consumption and investment habits.

2 Literature review

Elasticity of taxable income

There is a large body of literature dedicated to quantifying the elasticity of taxable income with respect to marginal income tax rates (Saez, Slemrod, and Giertz (2012) provide an excellent review). These papers generally find a significant estimated elasticity of taxable income. For example, Auerbach and Poterba (1988) find that changes in the capital gains tax rate induce large changes in capital gains realizations to avoid additional taxation. Additionally, Auerbach and Slemrod (1997) find that reducing the top individual tax rate to below the corporate tax rate in 1988 caused a shift in business activity towards pass-through entities, which are not subject to corporate tax. The consensus is that consumers are highly sensitive to tax considerations.

Behavioral responses to other taxes

Corporate taxes have also been shown to exert a strong influence on the investment behavior of affected entities. Hall and Jorgenson (1967) conduct one of the first studies of this nature, revealing that post-war tax credits for equipment expenditures were directly responsible for subsequent increases in business investment. There are numerous modern papers that further discuss the extent to which government policies can incentivize investment, including tax holidays in Latin America (Klemm and Van Parys 2012) and direct tax incentives in Sub-Saharan Africa (Van Parys and James 2010). There is also a growing body of research refuting the notion of corporations' elasticity of investment; Yagan (2015), for example, finds that the 2003 Dividend Tax Cut in the US "caused zero change in corporate investment and employee compensation," despite expectations to the contrary. These studies should provide a loose indication of how analogous individual income tax changes will affect private investment habits.

There are also numerous studies of the influence of other taxes (e.g., sales tax) on budget-sensitive individual behavior. Baker, Johnson, and Kueng (2017) provide recent evidence that taxpayers respond strongly to sales tax shocks, stockpiling items before tax increases and reducing subsequent spending even when the after-tax cost of goods has not changed. Many consumers, in fact, are so psychologically opposed to taxation that they will exhibit "a desire to avoid taxes ... that exceeds the rational economic motivation to avoid a monetary cost" (Sussman and Olivola 2011). Baker, Johnson, and Kueng (2017) also provide an excellent survey of papers that study similar policies. On a larger governmental scale, there have been numerous recent papers investigating the extent to which changes in national Value Added Taxes (VATs) influence consumer spending (Crossley, Low, and Sleeman (2014); Cashin and Unayama (2016); D'Acunto, Hoang, and Weber (2016)). However, because VATs are generally included in purchase prices and US sales taxes are not, the results may not be applicable to US consumers. Further study of US states' nebulous and often hidden sales taxes reveals "substantial deviations from frictionless optimizing behavior" (Chetty, Looney, and Kroft (2009); Finkelstein (2009); Cabral and Hoxby (2012)). Thus it is clear that increases in individual income tax rates should influence the budget-sensitive investment behavior of individual investors, but the extent and rationality of the change cannot be easily intuited.

My paper fills a relevant gap in the literature: there are few studies that discuss the response of individual-level tax-advantaged savings to increases in marginal income tax rates. My results will help illuminate the decision-making process of individual investors and serve as a guide for governments looking to understand the potential impact of increasing their income tax rates.

3 Conceptual framework

In this section, I set up a stylized two-period model of savings behavior to contextualize and structure my empirical analysis. I follow Chetty et al. (2014) and Moser and Silva (2019), who leverage similar pared-down models of savings for descriptive purposes.

Setup

An individual lives for two periods: $t = 1$ is her "working" period, and $t = 2$ is her "retirement" period. The government imposes a linear income tax $T \in (0, 1)$ in the first period to pay for public goods and retirement transfers.

At time $t = 1$, the individual earns a pre-tax wage W. She then decides how much of this income to save for retirement vs. consume in the first period. To reflect the tax advantage of saving for retirement, I allow retirement savings to be untaxed. Letting S and C_1 denote first-period savings and consumption respectively, I get the following budget constraint:

$$(1 - T)(W - S) = C_1. \tag{1}$$

At time $t = 2$, the individual is able to consume her retirement savings plus interest r. This results in total second-period consumption $C_2 = (1 + r)S$. I also allow for standard time discounting $\delta \in (0, 1)$ in addition to present bias $\beta \in (0, 1]$.

Thus her total utility can be expressed as follows, assuming a smooth, concave utility function $u(\cdot)$:

$$U(S) = u(C_1) + \beta \cdot \delta \cdot u((1 + r)S). \tag{2}$$

The individual chooses how much to save in the first period, maximizing utility subject to her budget constraint. Thus her problem is as follows:

$$\begin{aligned} \max_{S} \quad & U(S) = u(C_1) + \beta \cdot \delta \cdot u((1 + r)S) \\ \text{s.t.} \quad & (1 - T)(W - S) = C_1. \end{aligned} \tag{3}$$

The planner's goal is to choose a tax rate T that non-myopically maximizes welfare. If I let $g(\cdot)$ represent the concave utility from a public good and set $V(S)$ equal to unbiased utility (that is, $V(S) = u(C_1) + \delta \cdot u((1+r)S)$), then the planner's problem is as follows:

$$\begin{aligned} \max_{T} \quad & V(S) + g(T(W-S)) \\ \text{s.t.} \quad & (1-T)(W-S) = C_1. \end{aligned} \tag{4}$$

Discussion

Because $\delta < 1$, the consumer will always value future consumption less than current consumption. However, when $\beta < 1$ (indicating present bias), the consumer suboptimally under-weights future consumption and over-consumes in the first period. The government wants to remedy this imbalance by increasing T to further incentivize saving, and thus increase retirement consumption back to its optimal value. However, if $\beta \cdot \delta$ is sufficiently small, a large increase in T may cause the income effect to dominate the substitution effect, thus decreasing retirement savings.

In this paper, I study increases in marginal income tax rates that may implicitly make it more beneficial to save for retirement. In my model, these tax increases are akin to augmenting T with the goal of counteracting the efficiency loss induced by present bias. If the substitution effect outweighs the income effect, I expect my empirical analysis to show that increases in marginal income tax rates are associated with increased retirement investing in tax-deferred pension plans.

4 Data

IRS Form 5500

The goal of my study is to deduce the effect of marginal income tax rate changes on relevant retirement preparedness metrics across US states. To track my outcomes of interest I use the Form 5500, a mandatory annual IRS filing that details information about the participants, investments, and management of each private pension plan in the United States. The data are identified by employer identification number (EIN), employer name, plan name, plan number, and filing ID, allowing me to create a panel dataset that tracks each company's plans over time.

In order for this dataset to be relevant to my analysis, I must assume that each company's employees work and earn income in the same state that the company is registered in. Otherwise, income tax changes in the company's home state would have no conceivable influence on the investment habits of its employees. For this reason I restrict my analysis to plans with fewer than 100 participants, the IRS' own threshold for "large" vs. "small" companies. My assumption is that the employees of smaller businesses are more likely to work in the state that their company is registered in.

I use the primary Form 5500 for participant information and Schedule I for financial information. Starting in 2009 the IRS offered a one-document Form 5500-SF (Short Form) to small companies, and I leverage these filings as well. To select relevant dependent variables, I browse the IRS' official annotated Form 5500 to select fields of relevance. My outcome categories of interest, which I hypothesize may be affected by changes in marginal income tax rates, are: plan net assets; employer contribution amounts; participant contribution amounts; plan participation; and net income. In some cases I also consider the beginning-of-year (BOY) and end-of-year (EOY) values for a particular field (when available) to calculate changes for a given year. In addition to these outcomes of interest, I also record each plan's age and its 2-digit NAICS business code, an indicator of the type of work being conducted by each business. These covariates are then included as controls in most of my regressions.

I obtain a total of nearly 12 million relevant filings, which I then filter down to roughly 8.6 million. In addition to selecting only small plans, I also isolate to single-employer 401(k)-style defined contribution plans that exist during the relevant years of my analysis (1999-2018). I only consider each company's most recent filing for a given year, as many companies will submit incorrect filings and then revise them later. Finally, I aggregate these remaining filings into a panel dataset that contains one filing per plan between the years 1999 and 2018.

Due to the existence of outliers/typos in the original Form 5500 dataset, I Winsorize all numeric fields at the 99.5% level. Summary statistics for these Winsorized outcomes can be seen in Table 1. All my outcomes of interest have long right tails, as they are either bounded below by zero (e.g., participation and contributions) or contain values that are many multiples of the mean (e.g., plan net assets). The mean participant contribution amount is nearly twice the mean employer contribution amount, which is expected since most employers will only match a portion of their employees' contributions.

	Median	Min	Mean	Max	SD
Net assets ($)	435,375	0	1,068,842	17,886,812	1,898,684
Δ Net assets ($)	48,410	-3,378,657	76,304	2,968,283	428,362
Empl. contributions ($)	16,251	0	42,834	833,099	83,483
Partcp. contributions ($)	39,447	0	70,888	964,011	104,320
Total contributions ($)	75,278	-845,442	163,053	2,728,003	312,998
Net income ($)	47,996	-2,959,796	77,197	2,631,869	389,029
Total partcp.	16	0	28	326	36
Δ Total partcp.	0	-89	0	86	11
Partcp. w/ balance	12	0	21	228	27

Note: All outcomes are relatively small because the sample focuses on small businesses.

Table 1: Form 5500 summary statistics.

For computational ease, I then produce a further derived panel dataset that averages plans by year and state. This results in a 1,000-row dataset that is much more conducive to exploratory regressions than a nearly 9 million-row dataset.

Treatment

Because wealthier individuals exhibit much greater tax-evasive tendencies (Alstadsæter, Johannesen, and Zucman 2019), I track marginal income tax rate changes with respect to "high-earning" individuals. I do this in the following way. First, I select the 95th-percentile US income in 1999, which amounts to $142,000 (Bureau 1999). Then, I inflate this $142,000 income (using inflation numbers from the Bureau of Labor Statistics) through 2018 to produce 20 years of constant real annual incomes. Next, I run each year and its corresponding income through NBER's TAXSIM simulator (first outlined in Feenberg and Coutts (1993)) once per state, allowing me to capture the amount of income tax that would be paid for each state-year combination. I then report the marginal income tax rate from each TAXSIM simulation, resulting in a state-year panel of marginal income tax rates.

Because this dataset is non-binary (that is, each marginal tax rate is a real number), I devise a methodology to convert these continuous values into binary treatment. I begin by considering each state's year-on-year change in marginal income tax rate. I then isolate to changes that are "large" (magnitude greater than 0.25 percentage points), "isolated" (meaning no other changes in +/- two years), and positive. Finally, I remove data from the two states (Connecticut and Maine) where treatment occurs more than once. This results in a clean state-year panel dataset with eight instances of treatment: Arizona, Delaware, Illinois, Maryland, North Carolina, Pennsylvania, Vermont, and Wisconsin. The earliest treatment occurs in 2001, while the latest occurs in 2015.

Macroeconomic controls

In addition to Form 5500 data, I control for macroeconomic characteristics in each state and year that may have influenced my outcomes of interest, drawing from Bureau of Economic Analysis data. I account for state GDP, income, and high school enrollment, which Kuzniak and Grable (2017) suggest may influence retirement decision-making. I also include population and employment, which could affect the aforementioned factors. This allows me to create a panel dataset of controls by state and year, some of which are linearly interpolated due to missingness.

Once all data aggregation is complete, I create a 960-row panel that contains control, treatment, and outcome variables for the 48 remaining US states and each year between 1999 and 2018. I use this dataset to conduct my main analyses.

Exploratory analysis

I begin by exploring the data with some basic cross-sectional regressions to determine whether higher marginal income tax rates are associated with higher or lower rates of investment and participation in retirement plans. For each outcome, I estimate a naive regression of the following form:

$$Y_{st} = \alpha + \gamma_s + \nu_t + \beta r_{st} + \epsilon_{st} \tag{5}$$

Where Y is the specified outcome; s indexes states; t indexes years; α is a constant term; γ_s are group-fixed effects for state s; ν_t are time-fixed effects in year t; r_{st} is the marginal income tax rate in state s at time t; and ϵ_{st} is a stochastic error term for state s in year t such that $\mathbb{E}[\epsilon_{st}] = 0$. Here, β is the coefficient of interest.

I perform this regression once for each of my nine outcomes of interest. In two of

Table 2: Exploratory regressions with Form 5500 data

	Net assets	Change in total partcp.
	(1)	(2)
Marginal income tax rate	20,032.630***	−0.047**
	(4,125.457)	(0.023)
State fixed effects	Yes	Yes
Year fixed effects	Yes	Yes
N	960	960
R^2	0.959	0.823
Adjusted R^2	0.956	0.810
Residual Std. Error (df = 892)	66,273.460	0.365
F Statistic (df = 67; 892)	309.416***	62.076***

Notes: Table contains coefficient estimates from naive regressions of retirement investment outcomes against marginal income tax rate, with state and year fixed effects.

the nine regressions, I find a significant coefficient on *Marginal income tax rate*; these regressions are outlined in Table 2. All remaining regressions can be found in Appendix A.a. While I would prefer to see significant coefficients for a few more exploratory regressions, even two significant coefficients is enough to warrant continued investigation, particularly the strong positive coefficient on plan net assets.

5 Estimation strategy

Difference-in-differences

My primary causal specification is a standard difference-in-differences model that takes advantage of my previous definition of treatment. For each outcome, I use the full state-year panel dataset to estimate an equation of the following form:

$$Y_{st} = \alpha + \gamma_s + \nu_t + \delta Post_{st} + \theta X_{st} + \epsilon_{st} \tag{6}$$

Where $Y_{st}, \alpha, \gamma_s, \nu_t$, and ϵ_{st} are defined as above; $Post_{st}$ is a dummy that equals one if state s has been treated in year t or earlier; X_{st} is a vector of control variables for state s and year t; and δ is the unbiased estimator of treatment effect (Wing, Simon, and Bello-Gomez 2018).

Unless otherwise specified, I use the same control covariates across all analyses in this paper: plan age (estimated via Form 5500 history) and the macroeconomic controls mentioned above. Note that, for the sake of space, I will not include these coefficients in regression tables.

Restricted difference-in-differences

My second treatment-focused specification is designed to tease out more granular treatment effects than equation (6). I begin by further restricting my dataset, removing all states except those that are treated and those that have no state income tax during my relevant timeframe (Alaska, Florida, Nevada, New Hampshire, South Dakota, Tennessee, Texas, Washington, and Wyoming). In theory, this will allow me to use the zero-tax states as a benchmark against which my treated states can be measured, as the changes in marginal income tax rates in the zero-tax states will reflect and account for the potential influence of federal policy change. For the treated states, I then remove data for all years except +/- two years from treatment; this will help me work around the multiple-treatment issues that could arise from my treatment definition. Next, I remove all data for the zero-tax states except in years where data from treated states are being considered. This reduces my state-year panel from 960 rows to 208. While my dataset is now much smaller, the de facto comparison I have set up should reveal a more acute treatment effect.

Difference-in-differences with shifts

My third specification is a robustness check. I consider a setting that is similar to equation (6), except I include lag and lead treatment coefficients instead of a binary post-treatment indicator. This will allow me to generate lag-and-lead coefficient plots, which will help determine whether or not the parallel trends assumption is upheld in my regressions. I use my initial difference-in-differences specification as a baseline instead of the restricted version; this is because the restricted difference-in-differences specification removes too many observations for a lags-and-leads regression to have meaningful significance. Thus I estimate an equation of the following form:

$$Y_{st} = \alpha + \gamma_s + \nu_t + \sum_{j=-6, j\neq 0}^{6} \delta_j \mathbb{I}_{jst} + \theta X_{st} + \epsilon_{st} \tag{7}$$

Where all coefficients are the same as before except $\mathbb{I}_{jst}$ is the indicator that state s in time t was treated exactly j years ago and δ_j are now the coefficients of interest. The coefficients $\delta_j : j < 0$ allow me to observe any pre-trends in my data, while $\delta_j : j > 0$ demonstrate the hypothetical time-granular treatment effects. If the parallel trends assumption is upheld for a specific outcome, we should expect to see a significant "trend change" around $j = 0$.

6 Results

Difference-in-differences

I find that treatment has a significant impact on five of the nine outcomes of interest. In particular, when controlling for state/year/business sector fixed effects, statewide macroeconomic controls, and plan-specific controls, an increase in marginal state income tax rate is associated with increases in: plan net assets; participant contribution amounts; total contribution amounts; total plan participation; and number of plan participants with an account balance. See Table 3 for these results and Appendix A.b for all regression results (including the four regressions with insignificant coefficients of interest).

Restricted difference-in-differences

I find that, with the additional sample restrictions presented above, only one of the five previously-significant outcomes (plan net assets) is significantly affected by treatment, and the magnitude is reduced. However, the post-treatment coefficients in Table 4 are nearly all the same sign and roughly the same magnitude as their corresponding coefficients from the original difference-in-differences specification. This indicates that the regression may be statistically underpowered, as it exhibits the same trend but with lower significance.

Table 3: Significant difference-in-differences regressions with Form 5500 data

	Dependent variable:				
	Net assets	Partcp. contributions	Total contributions	Total partcp.	Partcp. with balance
	(1)	(2)	(3)	(4)	(5)
Post-treatment	51,549.140***	843.559*	5,954.089*	0.407**	0.374***
	(8,973.716)	(507.835)	(3,346.091)	(0.170)	(0.123)
State fixed effects	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes
NAICS fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	960	960	960	960	960
R^2	0.982	0.956	0.977	0.930	0.957
Adjusted R^2	0.980	0.951	0.974	0.922	0.953
Residual Std. Error (df = 862)	44,092.720	2,495.269	16,441.150	0.834	0.604
F Statistic (df = 97; 862)	494.716***	192.465***	369.439***	117.879***	199.623***

Notes: Table contains coefficient estimates from difference-in-differences regressions of retirement investment outcomes against post-treatment, defined as a large year-on-year increase in the marginal state income tax rate. Controls include state-level estimates of employment, GDP, population, income, and education; Form 5500-based estimates of plan age; and state, year, and NAICS fixed effects. This table includes only regressions with significant coefficients on treatment.

Table 4: Previously-significant restricted difference-in-differences regressions with Form 5500 data

	Dependent variable:				
	Net assets	Partcp. contributions	Total contributions	Total partcp.	Partcp. with balance
	(1)	(2)	(3)	(4)	(5)
Post-treatment	30,106.630***	220.831	−1,511.979	0.282	0.144
	(11,285.550)	(640.406)	(5,792.315)	(0.241)	(0.180)
State fixed effects	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes
NAICS fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	220	220	220	220	220
R^2	0.993	0.980	0.981	0.970	0.980
Adjusted R^2	0.990	0.971	0.973	0.957	0.972
Residual Std. Error (df = 153)	31,367.540	1,779.971	16,099.410	0.671	0.501
F Statistic (df = 66; 153)	331.891***	112.659***	122.201***	74.945***	116.118***

Notes: Table contains coefficient estimates from difference-in-differences regressions of retirement investment outcomes against post-treatment, defined as a large year-on-year increase in the marginal state income tax rate. Controls include state-level estimates of employment, GDP, population, income, and education; Form 5500-based estimates of plan age; and state, year, and NAICS fixed effects. These regressions use the restricted dataset with sample and timeframe restrictions. This table includes only regressions with (formerly) significant coefficients on treatment.

See Appendix A.c for the remaining difference-in-differences regressions with the restricted sample.

7 Robustness check

Testing parallel trends assumption

To test the parallel trends assumption, I consider the coefficients δ_{-6} through δ_{-1} in equation (7). If these coefficients are equal to zero, that indicates the parallel trends assumption is upheld. One downside of my regression in equation (7), however, is that I am underpowered to properly estimate the coefficients of interest. In Tables 5 and 6, which show these lagged coefficients for all nine outcomes of interest, there are effectively no significant coefficient values for δ_{-6} through δ_{-1} – this would seem to imply that the parallel trends assumption is upheld. However, since the treatment coefficients for $t > 0$ are also insignificant, it's impossible to tell if the parallel trends assumption is truly upheld or if my regressions are just too underpowered to detect the existing parallel trend.

See, for example, Figure 1, which shows lags-and-leads plots for +/- 6 years of treatment coefficients regarding plan net assets and participants with account balance (both of which exhibit a significant relationship with treatment in Table 3). Visually, if one were to ignore the error bars, it looks like a meaningful "jump" occurs at time $t = 0$ in both cases. This matches the significant and positive treatment effect observed in Table 3.

If the parallel trends assumption were upheld, we would expect to see precise zero coefficients at times $t < 0$; and if treatment were significant, we would expect to see precise nonzero coefficients at times $t > 0$. Observing Figure 1, it seems that the parallel trends assumption is upheld, as all negative coefficients are not significantly different from zero. However, all future coefficients are also not significantly different from zero, seemingly invalidating the treatment effects found in Table 3. Thus I am led to believe that the lags-and-leads regressions are statistically underpowered, and more data are required to determine which coefficients are truly zero and which are nonzero but lacking in statistical power. See Appendix B for the remaining seven lags-and-leads plots, nearly all of which suffer from the same lack of statistical power.

Table 5: Lagged difference-in-differences regressions with Form 5500 data (a)

	Dependent variable:				
	Net assets	Change in net assets	Empl. contributions	Partcp. contributions	Total contributions
	(1)	(2)	(3)	(4)	(5)
Pre-treatment (6 years)	−45,096.710*	−1,693.283	−1,594.241	126.369	−3,810.908
	(25,067.720)	(11,514.250)	(1,104.568)	(1,460.713)	(9,662.159)
Pre-treatment (5 years)	−38,388.100	−10,991.690	222.572	578.133	−11,664.450
	(23,631.960)	(10,854.770)	(1,041.304)	(1,377.050)	(9,108.756)
Pre-treatment (4 years)	−31,786.920	−3,692.338	−729.112	197.025	−3,354.136
	(22,501.770)	(10,335.640)	(991.503)	(1,311.193)	(8,673.131)
Pre-treatment (3 years)	−34,491.440	−2,400.686	−526.536	34.892	−3,790.440
	(22,473.940)	(10,322.860)	(990.277)	(1,309.572)	(8,662.407)
Pre-treatment (2 years)	−30,448.220	1,965.409	−354.248	−463.103	−2,052.183
	(21,526.170)	(9,887.524)	(948.516)	(1,254.345)	(8,297.096)
Pre-treatment (1 year)	−20,655.510	5,291.512	−255.277	−174.378	4,513.885
	(21,536.190)	(9,892.123)	(948.957)	(1,254.928)	(8,300.955)
State fixed effects	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes
NAICS fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	960	960	960	960	960
R^2	0.984	0.973	0.962	0.958	0.977
Adjusted R^2	0.982	0.969	0.956	0.951	0.974
Residual Std. Error (df = 830)	42,681.160	19,604.550	1,880.675	2,487.060	16,451.120
F Statistic (df = 129; 830)	397.705***	231.972***	163.457***	145.971***	277.699***

Notes: Table contains coefficient estimates from difference-in-differences regressions of retirement investment outcomes against lag and lead treatment coefficients, where treatment is defined as a large year-on-year increase in the marginal state income tax rate. Specification includes binary coefficients for +/- 6 years from treatment. Controls include state-level estimates of employment, GDP, population, income, and education; Form 5500-based estimates of plan age; and state, year, and NAICS fixed effects.

Table 6: Lagged difference-in-differences regressions with Form 5500 data (b)

	Dependent variable:			
	Net income	Total partcp.	Change in total partcp.	Partcp. with balance
	(1)	(2)	(3)	(4)
Pre-treatment (6 years)	−2,643.068	−0.198	0.025	−0.458
	(11,331.990)	(0.496)	(0.207)	(0.356)
Pre-treatment (5 years)	−11,675.110	−0.203	0.129	−0.319
	(10,682.940)	(0.468)	(0.195)	(0.336)
Pre-treatment (4 years)	−5,134.623	−0.443	0.019	−0.346
	(10,172.030)	(0.445)	(0.186)	(0.320)
Pre-treatment (3 years)	−3,905.241	−0.494	0.098	−0.444
	(10,159.450)	(0.445)	(0.186)	(0.319)
Pre-treatment (2 years)	1,373.700	−0.420	−0.052	−0.318
	(9,731.010)	(0.426)	(0.178)	(0.306)
Pre-treatment (1 year)	5,142.768	−0.164	0.018	−0.069
	(9,735.536)	(0.426)	(0.178)	(0.306)
State fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
NAICS fixed effects	Yes	Yes	Yes	Yes
Observations	960	960	960	960
R^2	0.974	0.931	0.847	0.959
Adjusted R^2	0.970	0.920	0.823	0.952
Residual Std. Error (df = 830)	19,294.220	0.845	0.353	0.606
F Statistic (df = 129; 830)	239.312***	86.626***	35.539***	149.443***

Notes: Table contains coefficient estimates from difference-in-differences regressions of retirement investment outcomes against lag and lead treatment coefficients, where treatment is defined as a large year-on-year increase in the marginal state income tax rate. Specification includes binary coefficients for +/- 6 years from treatment. Controls include state-level estimates of employment, GDP, population, income, and education; Form 5500-based estimates of plan age; and state, year, and NAICS fixed effects.

Figure 1: Lags-and-leads regressions for plan net assets and total participants, both 99.5% Winsorized

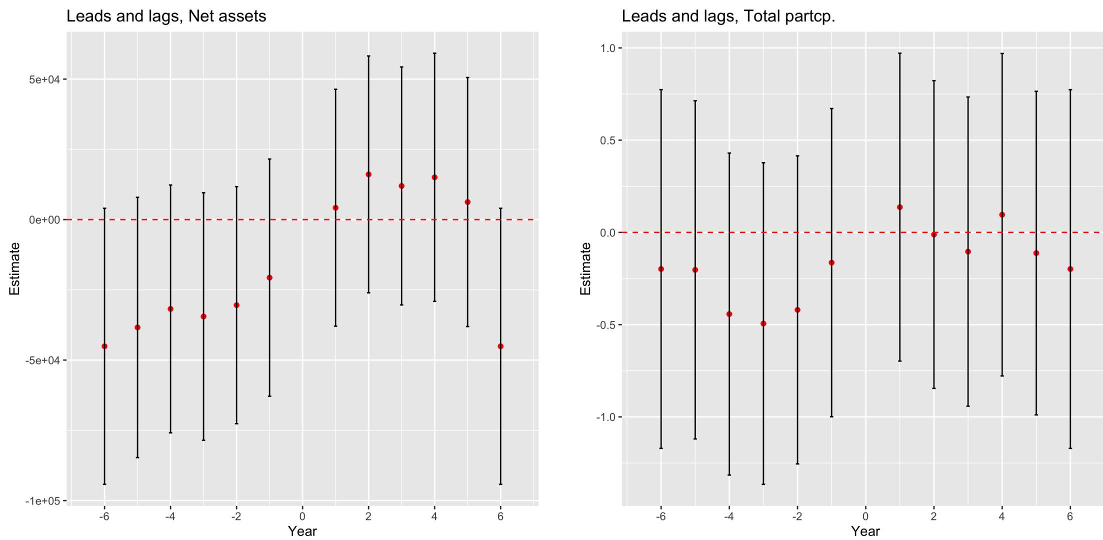

Note: Each figure shows 6 years of lag-and-lead treatment coefficients, which are estimated via a difference-in-differences regression of plan outcome against treatment dummies. Controls include state-level estimates of employment, GDP, population, income, and education; Form 5500-based estimates of plan age; and state, year, and NAICS fixed effects.

For the sake of argumentation, I will assume in the remainder of this paper that the parallel trends assumption is indeed upheld for my outcomes of interest and thus the difference-in-differences results are valid.

8 Discussion

Interpreting results

My results provide preliminary confirmation of the hypothesis that exogenous increases in marginal state income tax rate are associated with higher rates of saving and greater participation in tax-advantaged retirement plans. While my Form 5500-based dataset includes employees of all socioeconomic classes, I effectively isolate the effect on wealthier employees by only including legislation that has an outsize effect on high earners.

In my initial analysis following equation (6), I find that treatment is positively associated with plan net assets, contributions, and participation. My results of greatest interest are that treatment results in an $843 increase in plan participant contributions and a 0.37-person increase in number of participants with account balance. These are the two most important metrics for my analysis: if participants contribute more to their accounts or plans gain more active participants, this is a strong indication of increased retirement preparedness among employees.

My other significant results confirm this same trend, but the metrics themselves are less intuitive. First, I find that treatment is associated with a $51,549 increase in plan net assets. This result provides a sort of "robustness check" that I am not somehow failing to account for money leaving the plans, but it is less useful beyond that because plan net assets are composed of many components, some of which are not of interest in my analysis. I also find that treatment is associated with a 0.41-person increase in plan total participation. Since participants with account balance are a subset of total participants, it seems that there is only a 0.04-person increase in participants without account balance. This makes sense in the context of my analysis; when investing becomes more beneficial, one would expect those who open accounts to also invest in them. Finally, I find that treatment is associated with a $5,954 increase in total contribution amount. This result is somewhat mysterious because treatment is not associated with employer contribution amount, meaning roughly $5,000 of increased contributions are stemming from the "Other" category on the Form 5500. I will need to dig deeper to understand where this money comes from and what this result may indicate.

It is also worth mentioning that, although my lags-and-leads analysis contains no significant coefficients of interest, the insignificant coefficients are of the same sign and roughly the same magnitude as the significant coefficients resulting from equation (6). This indicates that, given more statistical power, my results in the lags-and-leads analysis may corroborate my results from the original regression specification.

Weaknesses

Data issues

In the Form 5500 dataset, there are many outliers that can be difficult to distinguish. For example, if a plan reports an abnormal negative contribution of large magnitude, that could either represent a typo or it could indicate large voluntary withdrawals that year; there's no way of knowing. Winsorization helps alleviate this issue at the extremes, but there is still the possibility that remaining errors marred my analysis. A separate issue with the Form 5500 data is that I depend on the assumption that employees work in the same state that their company is registered in, but I do not have sufficient data to prove this assumption.

A weakness with my treatment interpretation is that the cutoff of 0.25 percentage points is arbitrary, and there is no clear way to identify a non-arbitrary cutoff. Another issue is that, for some states, there is minor but constant fluctuation in marginal state income tax rate both before and after treatment, which may pose issues for analysis. I am unable to devise a metric for fluctuation that results in a rational threshold for exclusion, so I ultimately decide not to remove any states on the basis of pre- or post-treatment fluctuation.

An issue with my macroeconomic controls dataset is that widespread educational attainment data are not available for the entire period of interest, so I am forced to use high school enrollment as a proxy. This may not be an entirely accurate representation of my desired quantity of interest. I also include only two fields from the Form 5500 (NAICS code and plan age) as controls in my regressions, and deeper study of the Form 5500 itself may reveal other useful control covariates.

Attrition

Because only active companies are required to submit an IRS Form 5500, my primary panel dataset contains significant attrition and late entry. Any time a company dissolves between 1999 and 2018, that is considered "attrition" in my dataset, and whenever a company is created that qualifies as late entry. This is only an issue, however, if attrition and/or entry are correlated with treatment. Intuitively, it is hard to imagine a firm opening or closing as the result of a modest one-quarter-point increase in marginal income tax rate. However, since I am unable to test this hypothesis, it could be that my treatment is correlated with another covariate that influences attrition or entry in the dataset. Thus attrition and/or late entry cannot be conclusively ruled out as confounding factors in my analysis.

Lack of statistical power

As mentioned in the Results section, equation (7) produces an underpowered analysis. This prevents me from effectively testing the parallel trends assumption, and also means I am unable to estimate treatment at more granular time intervals. This does not necessarily invalidate my primary analysis, but it also fails to support it.

Future analysis

I would first like to devise an approach that tests granular year-specific treatment effects and the parallel trends assumption without being statistically underpowered. This would likely involve leveraging cloud computing resources to perform regressions on my entire disaggregated Form 5500 dataset. I would also like to locate a dataset that allows me to analyze the percent of employees that work in the state their company is registered in as a function of firm size. This would allow me to devise a more data-driven cutoff for firm size than simply "100." Finally, I would like to experiment with machine learning causal inference approaches to determine if there are more nuanced underlying trends in my data. Difference-in-differences is an excellent approach from an econometric perspective, but like all regression-based analyses it fails to account for non-linear trends in covariates of interest.

On a larger scale, I would like to expand my analysis to treatment interpretations with different base incomes. I would expect an ideal analysis to reveal that treatment effects are monotonically increasing in the reference income used to calculate marginal state income tax rate changes. In other words, increases in marginal income tax rate for impoverished Americans would not be associated with increased investment because these citizens lack the budget for any pre-tax savings, whereas wealthier employees with flexible investment habits would respond more elastically.

I would also like to include in my analysis a dataset containing information on non-pension private investment. This would allow me to better understand the investment actions of the employees I am studying. If, for example, treatment is associated with increased tax-deferred investment but decreased post-tax investment, that would indicate increased marginal income tax rates merely shift investment to a different vehicle as opposed to increasing the total dollar amount of savings.

Finally, I would like to include an additional individual-level dataset in my analysis tracking the same outcomes of interest. The Form 5500 is excellent because it is a mandatory public filing and thus conducive to panel dataset creation, but it lacks the ability to make any inference on an individual level. A more granular investment dataset (e.g., the Panel Survey on Income Dynamics) could allow me to perform analyses that study the same outcomes of interest but incorporate individual-level controls as well. This would also act as a secondary robustness check to confirm that my results are applicable to different methods of tracking the same metrics.

9 Conclusion

My paper is one of the first to provide preliminary evidence that an increase in marginal income tax rate is associated with improved metrics of retirement preparedness, including plan contributions and participation. My results indicate that increasing income tax rates may indirectly incentivize employees to save money for retirement. However, a more rigorous analysis of the parallel trends assumption is required before interpreting my results as true causality.

References

Alstadsæter, Annette, Niels Johannesen, and Gabriel Zucman (June 2019). “Tax Evasion and Inequality”. en. In: *American Economic Review* 109.6, pp. 2073–2103.

Andrews, Emily S. (1989). *Pension policy and small employers: At what price coverage?* English. An EBRI-ERF Policy Study Washington, D.C.: Employee Benefit Research Institute.

Auerbach, Alan J. and James Poterba (1988). “Capital Gains Taxation in the United States: Realizations, Revenue, and Rhetoric”. In: *Brookings Papers on Economic Activity* 1988.2, pp. 595–637.

Auerbach, Alan J. and Joel Slemrod (1997). “The Economic Effects of the Tax Reform Act of 1986”. In: *Journal of Economic Literature* 35.2. Publisher: American Economic Association, pp. 589–632.

Baker, Scott R., Stephanie Johnson, and Lorenz Kueng (Aug. 2017). *Shopping for Lower Sales Tax Rates*. en. Tech. rep. w23665. National Bureau of Economic Research.

Beshears, John et al. (2015). “"The Effect of Providing Peer Information on Retirement Savings Decisions"”. In: *The Journal of Finance* 70.3, pp. 1161–1201.

Beshears, John et al. (2021). “"Using fresh starts to nudge increased retirement savings"”. In: *"Organizational Behavior and Human Decision Processes"* 167, pp. 72–87.

Beshears, John et al. (2022). “"Present Bias Causes and Then Dissipates Auto-Enrollment Savings Effects"”. In: *"AEA Papers and Proceedings"* 112, pp. 136–141.

Bureau, US Census (1999). *Household Income: 1999*. EN-US. Section: Government.

Cabral, Marika and Caroline Hoxby (Nov. 2012). *The Hated Property Tax: Salience, Tax Rates, and Tax Revolts*. en. Tech. rep. w18514. National Bureau of Economic Research.

Cashin, David and Takashi Unayama (2016). “Measuring Intertemporal Substitution in Consumption: Evidence from a VAT Increase in Japan”. In: *The Review of Economics and Statistics* 98.2. Publisher: MIT Press, pp. 285–297.

Chetty, Raj, Adam Looney, and Kory Kroft (Sept. 2009). “Salience and Taxation: Theory and Evidence”. en. In: *American Economic Review* 99.4, pp. 1145–1177.

Chetty, Raj et al. (2014). “Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark”. English. In: *Quarterly Journal of Economics* 129.3, pp. 1141–1219.

Crossley, Thomas, Hamish Low, and Cath Sleeman (July 2014). *Using a temporary indirect tax cut as a fiscal stimulus: evidence from the UK*. en. Tech. rep. W14/16. Publication Title: IFS Working Papers. Institute for Fiscal Studies.

D’Acunto, Francesco, Daniel Hoang, and Michael Weber (Aug. 2016). *The Effect of Unconventional Fiscal Policy on Consumption Expenditure*. en. Tech. rep. w22563. National Bureau of Economic Research.

Feenberg, Daniel and Elisabeth Coutts (1993). “An introduction to the TAXSIM model”. In: *Journal of Policy Analysis and Management* 12.1. Publisher: John Wiley & Sons, Ltd., pp. 189–194.

Finkelstein, Amy (Aug. 2009). “E-ztax: Tax Salience and Tax Rates*”. In: *The Quarterly Journal of Economics* 124.3, pp. 969–1010.

GAO, U. S. (June 2015). “Retirement Security: Most Households Approaching Retirement Have Low Savings”. In: GAO-15-419.

Goss, Stephen C. (2010). “The Future Financial Status of the Social Security Program”. en. In: *Social Security Bulletin* 70.3.

Hall, Robert E. and Dale W. Jorgenson (1967). “Tax Policy and Investment Behavior”. In: *The American Economic Review* 57.3. Publisher: American Economic Association, pp. 391–414.

Klemm, Alexander and Stefan Van Parys (June 2012). “Empirical evidence on the effects of tax incentives”. en. In: *International Tax and Public Finance* 19.3, pp. 393–423.

Kuzniak, Stephen and John E. Grable (2017). “Does Financial Risk Tolerance Change over Time? A Test of the Role Macroeconomic, Biopsychosocial and Environmental, and Social Support Factors Play in Shaping Changes in Risk Attitudes”. English. In: *Financial Services Review* 26.4, pp. 315–338.

Moser, Christian and Pedro Olea de Souza e Silva (2019). “"Optimal Paternalistic Savings Policies"”. In: *"Columbia Business School Research Paper No. 17-51"*.

Saez, Emmanuel, Joel Slemrod, and Seth H. Giertz (2012). “"The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review"”. In: *Journal of Economic Literature* 50.1, pp. 3–50.

Soto, Mauricio (2008). “The effects of pension funding rules on the behavior of firms”. English. Ph.D. United States – Massachusetts: Boston College.

Sussman, Abigail B. and Christopher Y. Olivola (Feb. 2011). “Axe the Tax: Taxes are Disliked More than Equivalent Costs”. en. In: *Journal of Marketing Research* 48.SPL. Publisher: SAGE Publications Inc, S91–S101.

Van Parys, Stefan and Sebastian James (Aug. 2010). “The effectiveness of tax incentives in attracting investment: panel data evidence from the CFA Franc zone”. en. In: *International Tax and Public Finance* 17.4, pp. 400–429.

Wing, Coady, Kosali Simon, and Ricardo A. Bello-Gomez (2018). “Designing Difference in Difference Studies: Best Practices for Public Health Policy Research”. eng. In: *Annual Review of Public Health* 39, pp. 453–469.

Yagan, Danny (Dec. 2015). “Capital Tax Reform and the Real Economy: The Effects of the 2003 Dividend Tax Cut”. en. In: *American Economic Review* 105.12, pp. 3531–3563.

Appendices

A Regressions

A.a Exploratory regressions

Table 7: Exploratory regressions with Form 5500 data (a)

	Net assets	Change in net assets	Empl. contributions	Partcp. contributions	Total contributions
	(1)	(2)	(3)	(4)	(5)
Marginal income tax rate	20,032.630***	−212.041	199.791	−169.986	1,663.681
	(4,125.457)	(1,261.431)	(162.426)	(227.514)	(1,115.462)
State fixed effects	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes
N	960	960	960	960	960
R^2	0.959	0.969	0.922	0.902	0.971
Adjusted R^2	0.956	0.967	0.916	0.895	0.969
Residual Std. Error (df = 892)	66,273.460	20,264.280	2,609.290	3,654.902	17,919.350
F Statistic (df = 67; 892)	309.416***	416.307***	156.616***	122.560***	447.774***

Notes: Table contains coefficient estimates from naive regressions of retirement investment outcomes against marginal income tax rate, with state and year fixed effects.

Table 8: Exploratory regressions with Form 5500 data (b)

	Net income	Total partcp.	Change in total partcp.	Partcp. with balance
	(1)	(2)	(3)	(4)
Marginal income tax rate	42.498	0.042	−0.047**	0.071
	(1,244.075)	(0.082)	(0.023)	(0.062)
State fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
N	960	960	960	960
R^2	0.970	0.821	0.823	0.882
Adjusted R^2	0.968	0.808	0.810	0.873
Residual Std. Error (df = 892)	19,985.460	1.309	0.365	0.988
F Statistic (df = 67; 892)	427.676***	61.236***	62.076***	99.644***

Notes: Table contains coefficient estimates from naive regressions of retirement investment outcomes against marginal income tax rate, with state and year fixed effects.

A.b Difference-in-differences regressions

Table 9: Difference-in-differences regressions with Form 5500 data (a)

	Dependent variable:				
	Net assets	Change in net assets	Empl. contributions	Partcp. contributions	Total contributions
	(1)	(2)	(3)	(4)	(5)
Post-treatment	51,549.140***	2,731.917	74.673	843.559*	5,954.089*
	(8,973.716)	(3,986.078)	(380.210)	(507.835)	(3,346.091)
State fixed effects	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes
NAICS fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	960	960	960	960	960
R^2	0.982	0.972	0.961	0.956	0.977
Adjusted R^2	0.980	0.969	0.957	0.951	0.974
Residual Std. Error (df = 862)	44,092.720	19,585.750	1,868.179	2,495.269	16,441.150
F Statistic (df = 97; 862)	494.716***	308.778***	220.084***	192.465***	369.439***

Notes: Table contains coefficient estimates from difference-in-differences regressions of retirement investment outcomes against post-treatment, defined as a large year-on-year increase in the marginal state income tax rate. Controls include state-level estimates of employment, GDP, population, income, and education; Form 5500-based estimates of plan age; and state, year, and NAICS fixed effects.

Table 10: Difference-in-differences regressions with Form 5500 data (b)

	Dependent variable:			
	Net income	Total partcp.	Change in total partcp.	Partcp. with balance
	(1)	(2)	(3)	(4)
Post-treatment	2,995.485	0.407**	−0.104	0.374***
	(3,922.651)	(0.170)	(0.071)	(0.123)
State fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
NAICS fixed effects	Yes	Yes	Yes	Yes
Observations	960	960	960	960
R^2	0.973	0.930	0.844	0.957
Adjusted R^2	0.970	0.922	0.826	0.953
Residual Std. Error (df = 862)	19,274.100	0.834	0.349	0.604
F Statistic (df = 97; 862)	318.613***	117.879***	48.092***	199.623***

Notes: Table contains coefficient estimates from difference-in-differences regressions of retirement investment outcomes against post-treatment, defined as a large year-on-year increase in the marginal state income tax rate. Controls include state-level estimates of employment, GDP, population, income, and education; Form 5500-based estimates of plan age; and state, year, and NAICS fixed effects.

A.c Restricted difference-in-differences regressions

Table 11: Restricted difference-in-differences regressions with Form 5500 data (a)

	Dependent variable:				
	Net assets	Change in net assets	Empl. contributions	Partcp. contributions	Total contributions
	(1)	(2)	(3)	(4)	(5)
Post-treatment	30,106.630***	−4,043.194	−165.047	220.831	−1,511.979
	(11,285.550)	(6,704.245)	(554.990)	(640.406)	(5,792.315)
State fixed effects	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes
NAICS fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	220	220	220	220	220
R^2	0.993	0.979	0.982	0.980	0.981
Adjusted R^2	0.990	0.970	0.974	0.971	0.973
Residual Std. Error (df = 153)	31,367.540	18,634.060	1,542.562	1,779.971	16,099.410
F Statistic (df = 66; 153)	331.891***	106.655***	125.284***	112.659***	122.201***

Notes: Table contains coefficient estimates from difference-in-differences regressions of retirement investment outcomes against post-treatment, defined as a large year-on-year increase in the marginal state income tax rate. Controls include state-level estimates of employment, GDP, population, income, and education; Form 5500-based estimates of plan age; and state, year, and NAICS fixed effects. These regressions use the restricted dataset with sample and timeframe restrictions.

Table 12: Restricted difference-in-differences regressions with Form 5500 data (b)

	Dependent variable:			
	Net income	Total partcp.	Change in total partcp.	Partcp. with balance
	(1)	(2)	(3)	(4)
Post-treatment	−4,743.767	0.282	−0.087	0.144
	(6,799.151)	(0.241)	(0.123)	(0.180)
State fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
NAICS fixed effects	Yes	Yes	Yes	Yes
Observations	220	220	220	220
R^2	0.978	0.970	0.901	0.980
Adjusted R^2	0.969	0.957	0.859	0.972
Residual Std. Error (df = 153)	18,897.850	0.671	0.341	0.501
F Statistic (df = 66; 153)	104.151***	74.945***	21.155***	116.118***

Notes: Table contains coefficient estimates from difference-in-differences regressions of retirement investment outcomes against post-treatment, defined as a large year-on-year increase in the marginal state income tax rate. Controls include state-level estimates of employment, GDP, population, income, and education; Form 5500-based estimates of plan age; and state, year, and NAICS fixed effects. These regressions use the restricted dataset with sample and timeframe restrictions.

B Lags and leads figures

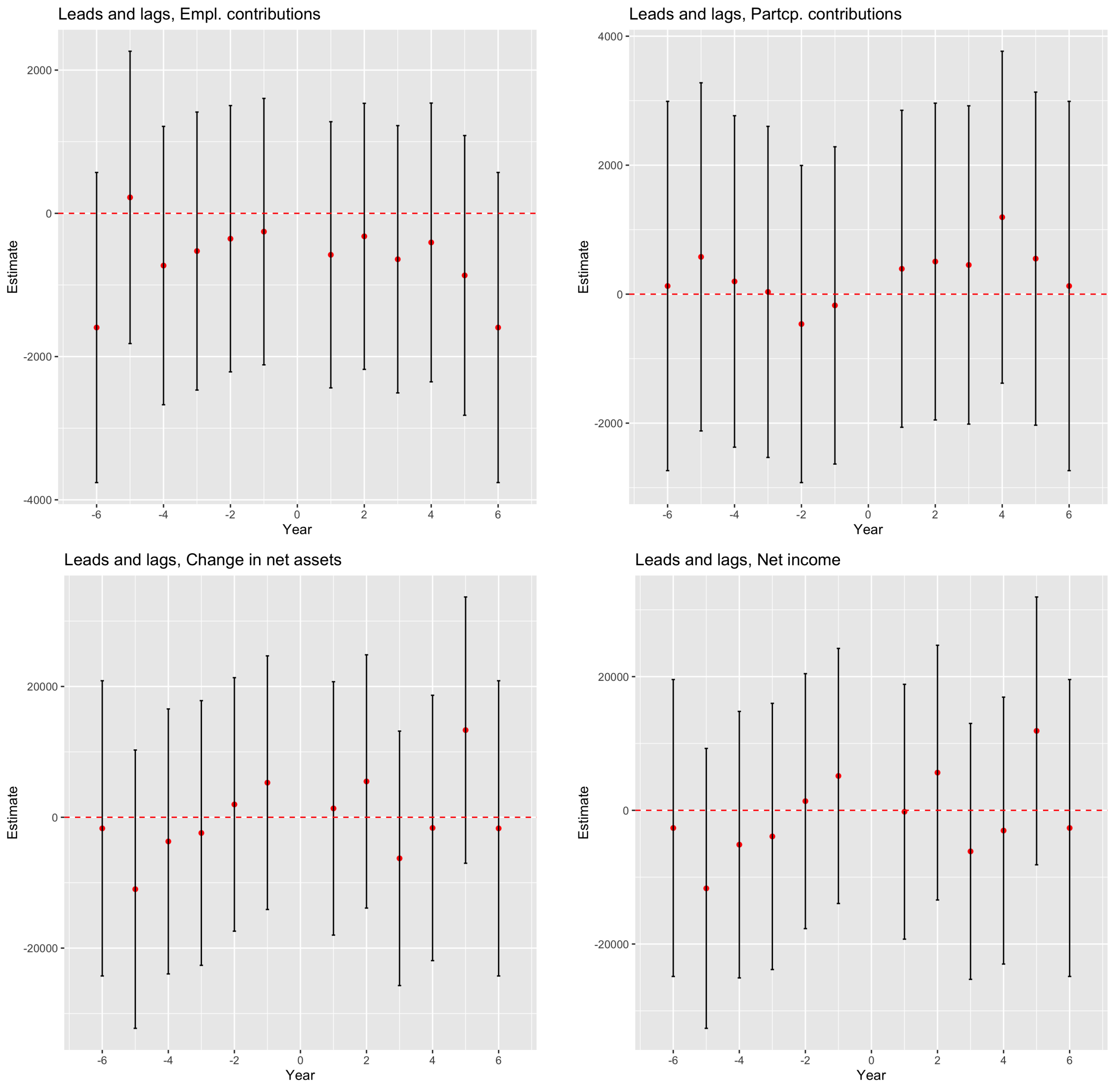

Figure 2: Remaining lags-and-leads regressions, all 99.5% Winsorized (a)

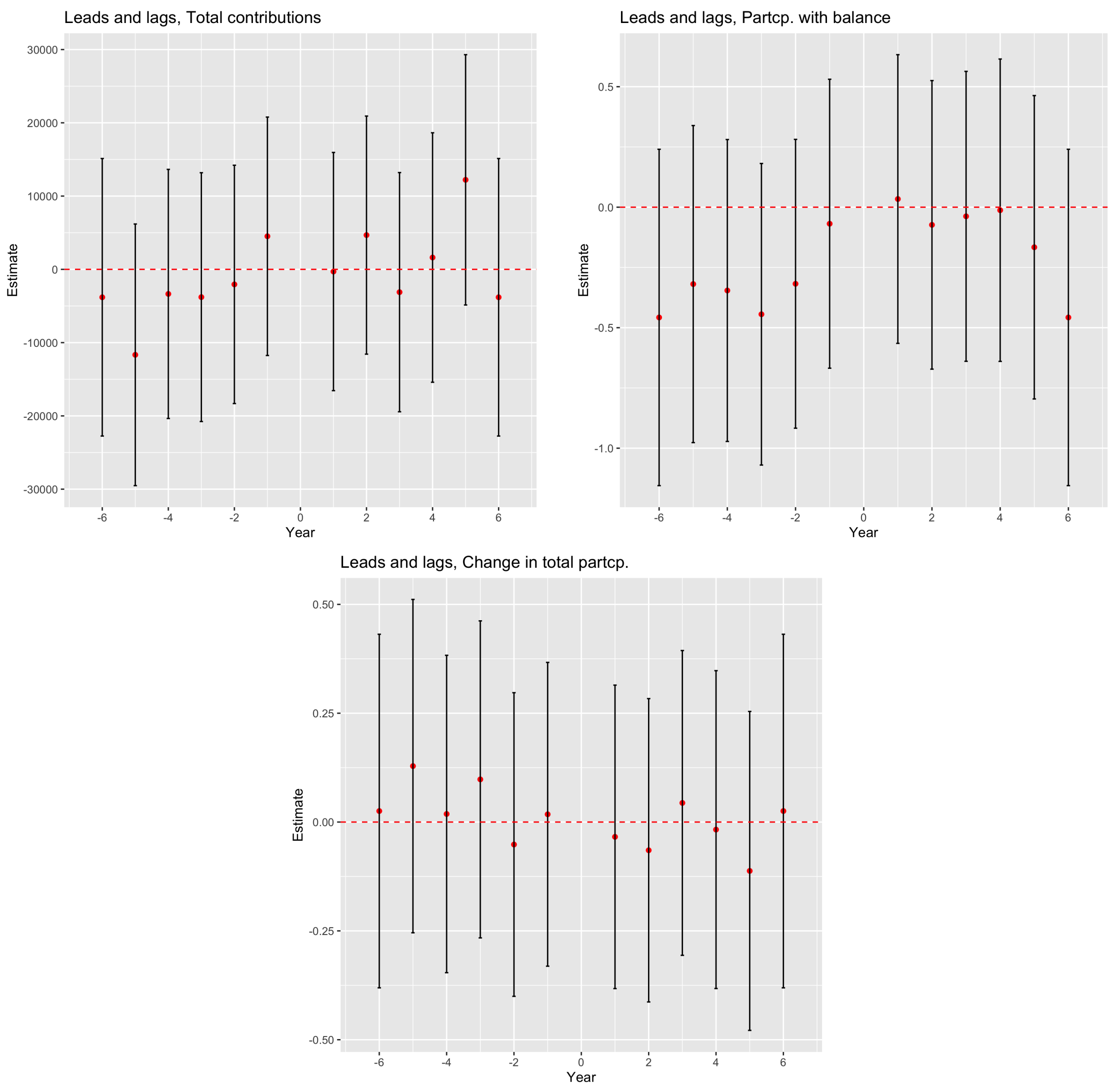

Figure 3: Remaining lags-and-leads regressions, all 99.5% Winsorized (b)