Does Raising Contribution Limits Lead to More Saving? Evidence from the 'Catch-up Limit' Reform

Adam M. Lavecchia*

October 23, 2014

Abstract

This paper studies the effect of raising contribution limits on retirement saving by exploiting the 'catch-up limit' provision, a rule which allows those over the age of 50 to make higher IRA and 401(k) contributions than those under the age of 50. Using an age-related regression discontinuity design, I recover an estimate of the causal effect of eligibility for 'catch-up limits' on saving, arguably permitting cleaner identification of the savings effect of these programs. I find that eligibility for 'catch-up limits' increases both average contributions by 20% and the probability of making a positive contribution by 3.5 percentage points (19%) for IRA owners, with no significant effects on 401(k) contributions. The average new contributions for those induced by the policy to contribute are relatively small. In the standard life-cycle model we expect that a contribution limit change will only affect agents for whom the limit is binding. However, these findings suggest that the initial response to eligibility for 'catch-up limits' is not limited to constrained individuals.

Keywords: Retirement saving; Tax-preferred savings accounts; Contribution limits; Regression discontinuity design

JEL: D14; H31; J26

*Email: adam.lavecchia@mail.utoronto.ca. University of Toronto, Department of Economics. 150 St. George Street, Toronto, Ontario, Canada M5S 3G7. I would like to thank my advisors, Michael Smart and Kory Kroft for their guidance throughout this project. Also, thanks to Michael Baker, Nicolas Gendron-Carrier, Michael Gilraine, Ashique Habib, Erica Lavecchia, Robert McMillan, Derek Messacar, and seminar participants at the Empirical Microeconomics (CEPA) seminar at the University of Toronto, the 48^{th} Annual Meeting of the Canadian Economics Association and EconCon 2014 for helpful comments and discussions. All remaining errors are my own.

1 Introduction

There is a growing concern that a significant number of households are not accumulating enough savings to smooth consumption in retirement.[1] Increasing life expectancies have exacerbated these concerns as current and future retirees risk out-living their savings. Furthermore, the decreasing availability of generous defined benefit pension plans implies that individuals will need to take a more active role in saving for retirement. Expanding the availability of IRAs and 401(k)s by increasing contribution limits is often cited as a way to increase personal saving.[2] However, the effectiveness of these programs has been widely debated in academic and policy circles with no clear consensus. One important issue in this debate is how expanding IRA and 401(k) limits will affect households who do not currently save in these plans. In this paper, I revisit this question using exogenous policy variation in IRA and 401(k) contribution limits.

Some of the first estimates on the effects of raising IRA and 401(k) contribution limits come from counterfactual policy simulations using structural models (Venti and Wise, 1990a, 1990b; Gale and Scholz, 1994). These studies estimate that raising contribution limits leads to a large (mechanical) increase in deductible contributions only for limit contributors, those whose constraint is binding. Using data from the Consumer Expenditure Survey (CES) and the Survey of Income and Program Participation (SIPP), Venti and Wise (1990a, 1990b) estimate that raising IRA limits leads to a 45 percent increase in annual contributions, more than two thirds of which represents new saving. However, using data from the Survey of Consumer Finances (SCF) and a model that allows tastes for saving to vary by IRA contributor status, Gale and Scholz (1994) estimate that only a small fraction of the increase in contributions following a limit change would represent new saving. These starkly different findings may be due to differences in econometric assumptions; the estimation in Venti and Wise (1990a, 1990b) assumes that differences in IRA saving across households is random after controlling for household characteristics, while Gale and Scholz compare IRA limit contributors, non-limit contributors and households that do not own IRAs.[3]

[1]See Poterba (2014) and the citations therein for one review of the state of retirement income security in the United States. Based on his reading of the literature, Poterba suggests that approximately 25 percent of households have not accumulated enough assets to purchase an annuity that would allow them to maintain even 50 percent of their pre-retirement consumption. He notes, however, that there is significant heterogeneity across households in retirement income security. In particular, replacement rates are relatively high for many low-income households given current expected Social Security benefits.

[2]Individual Retirement Accounts (IRAs) and employer-sponsored 401(k) savings plans are the two most popular tax-preferred savings accounts. Asset balances in IRAs, more than \$4.8 trillion in 2012, account for more than 25% of retirement wealth and 9.3% of household assets (Holden and Bass, 2014). 401(k) plans are the most popular employer-sponsored defined contribution plan. 401(k) balances totaled more than \$4.2 trillion in 2013, almost 50 percent of private pension assets in the United States (Holden and Schrass, 2014).

[3]In the Canadian setting, Milligan (2003) studies a reform that increased contribution limits for Registered Retirement Savings Plans (RRSPs). He finds that increases in *future* contribution limits lower *current* period contributions for those who are both unconstrained by the current period limit and expect to be constrained in the future. Milligan reconciles these findings with a model in which savers who expect their desired retirement account contributions to be

These mixed findings lead to the development of a large literature in public finance that estimates whether IRA and 401(k) contributions displace saving in taxable assets.[4] Comparing cross-sectional differences in financial assets between those employed by firms that offer a 401(k) and those who are not, Poterba, Venti and Wise (1995) conclude that the majority of contributions to 401(k)s represent new saving.[5] However, other empirical strategies that compare new versus existing IRA contributors (Attanasio and DeLeire, 2002), similar individuals by matching on a propensity score (Benjamin, 2003), 401(k) savers facing different employer match rates (Engelhardt and Kumar, 2007) and those at the borderline of eligibility for an IRA subsidy (Ramnath, 2013) find that the majority of contributions are financed by reducing saving in taxable accounts. In a recent study using within-person changes in 401(k) eligibility, Gelber (2011) finds evidence consistent with both large overall increases or decreases in saving due to 401(k) eligibility.[6]

The evidence in these papers has led to improvements in our understanding of who uses IRAs and 401(k)s, as well as on how the features of these plans affect savings behavior. However, the empirical strategies used are faced with important limitations (Bernheim, 2002). Historically, eligibility for IRAs and 401(k)s has either been universal or correlated with potential tastes for saving such as income or workplace pension status. Subsequently, cross-sectional comparisons of those eligible against those ineligible may confound the savings effect of these plans with differences in unobserved tastes for saving, even conditional on observable characteristics. Strategies that exploit within-person changes in IRA or 401(k) contributor status require that the time varying unobservable characteristics that influence saving evolve similarly for those whose contributor status switches and those whose does not. Significant differences in observables between these two groups may be indicative that unobservable tastes for saving also evolve differently over time.[7] Empirical evidence on the causal effects of raising contribution limits is especially scarce, in part because

constrained in the future shift some of their savings to the current period. In this case, increasing future contribution limits relaxes the future constraint, leading to lower contributions in the current period.

[4]In the Danish setting, Chetty et al. (2014) show that the overall savings effect of retirement accounts depends on whether individuals are active or passive in their savings allocation decisions. For those that are passive (about 85 percent of the sample), mandatory savings policies increase savings since individuals do not re-optimize by reducing saving in other accounts. Tax incentives to save in retirement accounts also have little effects on this group. For those that are active optimizers (about 15 percent of the sample), tax-incentives to save lead to reductions in other taxable saving, limiting the overall savings impact of these plans.

[5]Feenberg and Skinner (1989) use a similar empirical strategy, by comparing the taxable capital income of households who do and do not own an IRA and also find little evidence of substitution.

[6]A complementary large and growing literature in behavioral economics explores whether non-price mechanisms increase saving. Several papers have found that policies such as default options (Madrian and Shea, 2001; Choi et al., 2004) requiring active decisions (Carrol et al., 2009), salience (Duflo et al., 2006, 2007; Saez, 2009; Chetty et al., 2014), the actions of agents in an advisory capacity (Duflo et al., 2006), reminders (Karlan et al., 2010) and informational nudges (Clark et al., 2014) can be effective at increasing saving.

[7]For example, in Attanasio and DeLeire (2002), new IRA contributors are younger, have lower income, lower levels of financial assets and fewer children (on average) than continuing contributors. In Gelber (2011), employees newly eligible to participate in a 401(k) are younger, have lower income, less liquid and illiquid financial assets have have less debt (both secured and unsecured), on average, than employees who have always been eligible for their firm's 401(k) plan.

changes in statutory limits are infrequent. Interpreting evidence that compares IRA and 401(k) contributors before and after a limit change is difficult since changes in unobserved tastes for saving may be correlated with the limit change.

Using an empirical strategy that aims to address the limitations of prior research, this paper asks whether raising IRA and 401(k) contribution limits increases saving. The empirical strategy exploits the 'catch-up limit' provision, a rule introduced in 2002 which allows individuals over the age of 50 to make larger IRA and 401(k) contributions than those under the age of 50. The 'catch-up limit' provision admits an age-based regression discontinuity design, arguably allowing for cleaner identification of the savings effect of these programs. The validity of this empirical strategy requires only that all unobserved characteristics that influence saving are continuous in age at age 50. In Sections 4 and 5, I present evidence to support the plausibility of this identification assumption. In particular, using savings behavior from the years prior to the introduction of 'catch-up limits', I test whether IRA and 401(k) contributions are continuous at age 50. Reassuringly, I find no evidence that contributions or the probability of making a positive contribution are discontinuous at age 50.

Using pooled cross-sections from the SIPP for 2002, 2004 and 2005, I estimate that eligibility for 'catch-up limits' increases IRA contributions by \$78 (20%) and the probability of making a positive contribution by 3.5 percentage points (19%) for account owners. Moreover, average new contributions by those induced to participate are relatively small and eligibility for 'catch-up limits' has insignificant effects on 401(k) contributions.

What can explain these findings? As I show in Section 3, in the standard life-cycle model (LCM), we do not expect that a changing contribution limit will affect the choices of agents for whom the limit is not binding. Extending the standard model to include fixed costs of contributing to a retirement account can lead some agents to begin saving in the account. Specifically, some non-contributors who would have saved at the limit in the absence of the fixed cost may be induced to participate following a limit change. However, for participation to be optimal, the new contributions of these agents must be greater than the pre-reform limit; otherwise contributing under the old limit would have been optimal.

That eligibility for 'catch-up limits' leads to a large increase in the probability of contributing to an IRA and that the new contributions of those induced to participate are relatively small suggests that the initial response to the reform was not limited to those constrained by the limit. However, an increase in advertising by financial institutions directed towards those over the age of 50 (and eligible for 'catch-up limits') may be able to explain the large increase in IRA participation rates. This potential mechanism is supported by descriptive evidence from financial industry trade publications immediately following the reform. In these publications, industry experts advise bank executives to direct their advertising expenditures towards IRA owners over the age of 50 in order to increase assets under management. This advertising may serve as a reminder to individuals about the importance of saving for retirement or about the tax benefits associated with contributing to tax-advantaged accounts. Since only 28 percent of all IRA owners make a tax-deductible contribution each year, financial institutions have a large incentive to increase the number of active contributors. Differences between the effects of eligibility for 'catch-up limits' on saving in IRAs and 401(k)s, accounts only offered by employers and not financial institutions, also supports the plausibility of this mechanism. Though suggestive, this evidence is consistent with the findings from a recent but growing literature which shows that advertising has a significant effect on the demand for retail banking products (Gurun, et al., 2013; Hastings et al., 2013; Honka et al., 2014), and may speak to the need for policy makers to consider the possible supply-side effects of reforms that aim to increase retirement saving (Duarte and Hastings, 2012).

I then turn to the question of whether increases in IRA saving lead to lower saving in taxable accounts, such as checking and stock or mutual fund accounts. To control for the possibility that IRA saving is correlated with unobserved determinants of taxable saving, I use eligibility for 'catch-up limits' as an instrument for IRA saving. I find that increases in IRA contributions lead to modest but statistically insignificant increases in taxable saving. However, wide confidence intervals mean that I cannot also rule out modest crowd-out. For example, a $1 increase in IRA saving is predicted to increase saving in interest-bearing bank accounts by $0.44, though the lower bound of the 95% confidence interval suggests that modest crowd-out of up to $0.23 is possible.[8]

The remainder of this paper is organized as follows. Section 2 describes the 'catch-up limit' provision. Section 3 presents a stylized two period model that includes many of the features of the U.S. retirement savings system. I use the model as a benchmark to illustrate why the empirical findings are difficult to reconcile with the standard LCM. In Section 4, I describe the data and the empirical strategy used to identify the effect eligibility for 'catch-up limits' on IRA and 401(k) contributions. Section 5 reports the main results for IRAs and 401(k)s as well as several robustness checks, while Section 6 reports the crowd-out estimates. Section 7 provides some concluding remarks.

2 The 'Catch-up Limit' Provision

This section discusses the 'catch-up limit' provision, an IRA and 401(k) contribution rule introduced as part of the *Economic Growth and Tax Relief Reconciliation Act* (EGTRRA). In the empirical analysis below, I exploit this provision to recover an estimate of the causal effect of increasing contribution limits on retirement saving. Along with Social Security benefits, IRAs and 401(k)s

[8]Conversely, the upper bound of the 95% confidence interval can not rule out crowd-in of up to $1.11.

account for the vast majority of retirement income for most Americans.[9] Contributions to both of these accounts are tax-deductible in the year the contribution is made (up to an annual limit).[10] For 401(k)s, employee contributions are often supplemented by matching employer contributions, up to a maximum percentage of earnings or a fixed dollar amount (Benartzi and Thaler, 2007; Carrol et al., 2009). Interest earned on assets held in IRAs and 401(k)s accumulates tax-free. With some exceptions, withdrawals from these plans are subject to a 10% penalty for savers younger than 59.5 years of age.

The EGTRRA enacted several reforms ranging from reductions in marginal income tax rates to changes in estate and gift tax regulations. It also legislated changes to retirement savings plans, including IRAs and 401(k)s. For example, the EGTRRA introduced the Saver's Credit, a non-refundable tax credit for certain low-income savers. EGTRRA also allows individuals to transfer in-kind (roll-over) assets from one workplace pension to another 401(k) or IRA. The focus of this paper, however, is on the contribution limit increases that EGTRRA legislated over the 2000s. The IRA contribution limit was increased for the first time since 1981; from \$2,000 in 2001 to \$3,000 in 2002-2004, \$4,000 for 2005-2007 and \$5,000 from 2008 onwards.[11] Similar contribution limit increases were implemented for 401(k)s, though initial contribution limits were much higher.

The EGTRRA also introduced the 'catch-up limit' provision, a rule which allows individuals over the age of 50 to make larger IRA and 401(k) contributions than those under the age of 50. Beginning in 2002, individuals who turn 50 years old (or older) by the end of the calendar year are eligible for 'catch-up limits'; eligibility is therefore a deterministic function of an individual's year of birth. 'Catch-up limits' are intended to provide a way for older workers who previously had not saved sufficiently for retirement, an opportunity to "catch-up" (H.R. Rep. No. 107-51). For IRAs, the provision allowed eligible taxpayers to contribute an additional \$500 annually from 2002 to 2005 and \$1,000 annually from 2006 onwards.[12] Individuals are not permitted to carry-forward unused contribution limits to future years. Figure 1 plots the evolution of regular IRA contribution limits and 'catch-up limits'. This variation in the ability to contribute to an IRA or 401(k) is at the individual level, as opposed to the household level. For example, in 2002 joint tax

[9] See Poterba (2014) for one discussion of sources of retirement income for Americans. For a detailed description of Social Security Benefits, see http://www.ssa.gov.planners/about.htm.

[10] For 401(k)s the annual nominal dollar (employee + employer) contribution limit was \$45,465 in 1978, the first year this account was available. The Tax Reform Act of 1986 (TRA86) limited employee contributions to \$7,000 per year. This limit has been raised periodically and is \$17,500 for 2014. For IRAs, contributions were initially limited to \$1,500 in 1974. The Economic Recovery Act of 1981 expanded the availability of IRAs and increased the annual contribution limit to \$2,000. TRA86 limited the ability of individuals with high incomes and workplace pensions to deduct IRA contributions. For the 2014 calendar year, single households (tax-filers) with workplace pensions and a modified annual gross income (AGI) between \$60,000 and \$70,000 can partially deduct contributions. Tax-filers with a modified AGI above \$70,000 may make non-deductible contributions up to the annual limit.

[11] From 2008 onwards, EGTRRA mandated that IRA contribution limits be index to inflation and rounded to the nearest \$500 nominal dollar amount. For 2014, the contribution limit is \$5,500.

[12] The *Pension Protection Act* of 2006 made 'catch-up limits' (as well as the regular contribution limit increases in the EGTRRA) permanent.

Figure 1: The Evolution of IRA Contribution Limits

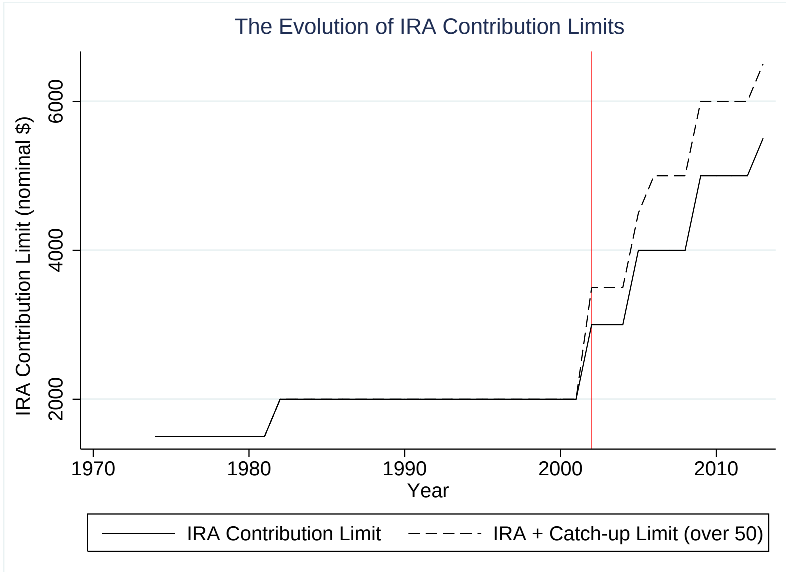

Notes: This figure plots nominal IRA contribution limits against time. The solid line represents the regular IRA contribution limit while the dashed line represents the regular IRA limit plus the 'catch-up limit' following the EGTRRA.

filers with both spouses over the age of 50 were eligible to contribute an additional $1,000 to IRAs ($500 for each spouse), while joint tax filers with one spouse over the age of 50 were only able to contribute an additional $500. Table 1 summarizes the main retirement income security provisions in the EGTRRA. Importantly, 'catch-up limits' are the only provision which differentially affect individuals around age 50.

Even before implementation, 'catch-up limits' were controversial. On April 2, 2001, the Center on Budget and Policy Priorities argued that the planned regular contribution limit increases and the 'catch-up limit' provision would not benefit the majority of Americans or lead to significant increases in retirement savings:

"The proposal would have virtually no effect, however, on families and individuals who do not make any deposits to IRAs under the current law or who deposit less than the current $2,000 limit. This proposal would directly benefit only those already making the $2,000 maximum contribution...Those at the limit almost certainly are among the most affluent of the taxpayers eligible for IRAs." (Orszag and Orszag, 2001)

This follows directly from the standard life-cycle model (LCM), which predicts that only constrained agents are affected by a limit change.

On the other hand, many in the media and the financial planning community praised the reform as a way to help procrastinators. For example, a March 2001 Sun-Sentinel newspaper[13] article describes the 'catch-up limit' provision in the following way: "These provisions would be a boon to those Baby Boomers who just turned 50 and haven't saved diddly, to stay-at-home moms re-entering the workforce, or to working parents who finally got their kids in college and need to attend to their own needs" (Cruz and Lade, 2001). The idea that eligibility for 'catch-up limits' could increase retirement account participation rates or increase in saving for those not constrained by the initial limit offers a competing prediction for the effect of this provision on IRA and 401(k) savings behavior.

Financial institutions also favored the the implementation of 'catch-up limits', viewing the provision as a way to encourage their customers to make additional IRA contributions (Dalton, 2001). Given the low marginal costs associated with administering these accounts, the expense fess charged for products such as mutual funds would likely make increasing IRA assets under management profitable. Some descriptive evidence suggests that financial institutions actively directed promotional activities towards those over the age of 50 in order to increase IRA deposits following EGTRRA. A 2004 article in the American Bankers Association's ABA Trust & Investments magazine encourages members to actively promote 'catch-up' contributions to clients as "part of a bank's overall strategy to increase assets under management" (Ellens, 2004). Another journal for executives in the banking sector recommends implementing specific advertising and promotional campaigns directed towards those eligible for 'catch-up limits': "For example, send a direct mail piece explaining the new 'catch-up' contributions to your IRA owners who are age 50 or older...Not only will you educate members about the IRA changes, you will create member awareness of the products and services your credit union offers." (Zuehlke, 2001). Other industry articles discuss "talking points" for bank tellers who recognize opportunities to recommend their institutions' IRA products to those eligible for 'catch-up limits' (Teller Sense, 2003). Overall, this evidence suggests an active effort by financial institutions to increase IRA contributions for those over the age of 50.

Poterba, Venti and Wise (1995) and Bernheim (2002) allude to the potentially important role of promotional activities for explaining aggregate IRA contributions. Poterba, Venti and Wise (1995) argue that the large decline in aggregate IRA contributions following the Tax Reform Act (TRA86) may be due to the decline in promotional activity by financial institutions. TRA86 limited the deductibility of IRA contributions for high-income households, leaving incentives unchanged for lower income households. The authors suggest that the large and persistent declines in IRA contributions beginning in 1987 for low- and middle-income households unaffected by the provisions in TRA86 may be due to a reduction in advertising by banks.

How might IRA promotional activity affect contributions? The marketing literature distinguishes between the "informative" and "persuasive" effects of advertising (Bagwell, 2007). This

[13] The Sun-Sentinel is the main newspaper for the Broward and South Plam Beach counties in Florida.

suggests that advertising by financial institutions may affect the demand for IRAs through one (or both) of these channels. Though research on the effects of bank advertising on retirement account contributions is limited, some recent studies suggest that both channels may be important in the demand for checking accounts and mortgages (Gurun et al., 2013; Hastings et al., 2013; Honka et al., 2014). In Section 5 I discuss some suggestive evidence that supports the plausibility of this mechanism.

Table 1: Summary of EGTRRA Pension and Retirement Income Security Reforms

Reform	Targeted Group
Personal marginal tax rate reductions	All
(Regular) IRA contribution limit ↑	All
'Catch-up Limits'	Age ≥ 50
DB pension contribution limit ↑	All DB pensions
DC plans (incl. 401(k)) limit ↑	All DC pensions
DC pension rule for top employees relaxed	All DC pensions
Saver's Tax Credit	AGI below cutoff
Credit for small firms starting pension plans	All eligible firms
Vesting of employer matches to DC plans relaxed	All DC pensions
DC withdrawals hardship rules relaxed	All DC pensions
Rollovers for certain DC pensions	All eligible DC pensions

3 Conceptual Framework

3.1 *Model Setup*

This section presents a stylized two-period, partial equilibrium model that illustrates the predictions of the standard life-cycle model (LCM) for an increase in contribution limits on retirement account participation rates, contributions and total personal saving[14] The model environment includes the institutional features of IRAs and 401(k)s discussed in the previous section.

Individuals, indexed by the subscript i, live for two periods, are endowed with exogenous income y in the first period and receive no income in the second period. Period 1 income can either be consumed or saved in one of two accounts: a retirement account or a taxable account. Retirement account contributions (R_{1i}) are tax-deductible so that taxable income in period 1 is $y - R_{1i}$. These contributions must be non-negative and cannot exceed an exogenous contribution limit $L > 0$. Individuals who contribute to the retirement account also pay a one-time fixed cost $k_i \geq 0$. For simplicity, I assume that k_i takes one of three values, k_H k_M, or k_L with $k_H > k_M > k_L = 0$. This

[14]The notation in this model is similar to Milligan (2003).

cost is independent of the contribution amount and may include the time cost of determining the asset allocation for a contribution or any fees charged by the financial institution administering the account. Individuals also face a constant marginal tax rate τ_1 in period 1.[15] Contributions to both the retirement account and the taxable account (N_{1i}) earn the same real pre-tax interest rate $r > 0$. In period 2, individuals face a constant marginal tax rate τ_2 and consume all after-tax resources. I assume that $0 < \tau_2 \leq \tau_1 < 1$ as most Americans face a higher marginal tax rate during their working years than in retirement. Both principal and interest earned on contributions to the retirement account are taxable in period 2. Interest earned on contributions to the taxable account is taxable in period 2 while principal is not. The period 2 budget constraint is $c_{2i} = (1+r)(1-\tau_2)R_{1i} + (1 + r(1-\tau_2))N_{1i}$. I also assume that $N_{1i} \geq 0$.

Given this set-up, contributions to the retirement account earn a higher after-tax real rate of return than taxable savings.[16] As a result, it is optimal for individuals to save in the taxable account only after contribution room in retirement account is exhausted. In practice, however, many individuals save in taxable accounts even though they contribute less than the limit in their IRA or 401(k). This may be due to a demand for liquidity since IRAs and 401(k)s have penalties for early withdrawals.[17] In the spirit of Chetty et al. (2014), I model the liquidity cost of saving in the retirement account with the convex function $v(R_{1i})$ with $v(0) \geq 0$, $v'(R_{1i}) > 0$ and $v''(R_{1i}) > 0$.

Preferences are represented by the utility function $U_i = u(c_{1i}, c_{2i}) - v(R_{1i})$ with $u_1 > 0$, $u_{11} < 0$, $u_2 > 0$, $u_{22} < 0$ and $u_{12} \geq 0$. Individuals choose a retirement account and a taxable account contribution to maximize lifetime utility subject to their period 1 budget constraint, period 2 budget constraint, the non-negativity constraints and the contribution limit constraint.

$$\max_{R_{1i}, N_{1i}} \; u(c_{1i}, c_{2i}) - v(R_{1i}) \tag{1}$$

subject to:

[15] The assumption of a constant marginal tax rate implicitly assumes that changes in R_{1i} do not cause the individual to move into a different tax bracket. Relaxing this assumption to allow marginal tax rates to vary with retirement account contributions does not affect the predictions discussed below.

[16] A \$1 pre-tax contribution to the retirement account costs $\$(1-\tau_1)$ in period 1 consumption and buys $\$(1+r)(1-\tau_2)$ in period 2 consumption. The sum of these two terms is $\tau_1 - \tau_2 + r(1-\tau_2)$. On the other hand, a \$1 pre-tax contribution to the taxable account costs \$1 in period 1 consumption and buys $\$1 + r(1-\tau_2)$ in period 2 consumption. Since $\tau_1 \geq \tau_2$, a \$1 contribution to the retirement account buys (weakly) more period 2 consumption than a \$1 contribution to the taxable account. In a model with more than two periods, the retirement account has the further advantage that interest accumulates tax-free while in the account.

[17] Gale and Scholz (1994) model this demand for liquidity with a three period structural model where income is stochastic. Faced with the possibility of a low draw of income in subsequent periods, individuals optimally save positive amounts in a taxable savings account even when contributions to the retirement account are less than the contribution limit.

$$c_{1i} = \begin{cases} (y - R_{1i})(1-\tau_1) - N_{1i} - k_i & \text{if } R_{1i} > 0 \\ y(1-\tau_1) - N_{1i} & \text{if } R_{1i} = 0 \end{cases}$$

$$c_{2i} = (1+r)(1-\tau_2)R_{1i} + (1 + r(1-\tau_2))N_{1i}$$

$$R_{1i} \geq 0;\ R_{1i} \leq L;\ N_{1i} \geq 0$$

The five possible optimal choices for R_{1i} and N_{1i} are described in the appendix. Let the pair (R^*_{1i}, N^*_{1i}) be the optimal choices for retirement and taxable account contributions for individual i. Intuitively, individuals weigh marginal contributions to the retirement account by comparing the additional benefit associated with higher retirement savings $u_2(\tau_1 - \tau_2 + r(1-\tau_2))$ against the marginal disutility from saving in an illiquid asset, $v'(R^*_{1i})$.

Due to the fixed cost of saving in the retirement account, positive contributions must also provide a higher level of utility than not contributing. Specifically, in addition to satisfying the relevant Kuhn-Tucker conditions, any solution for which positive retirement account contributions is optimal $(R^*_{1i} \in (0, L]\ , N^*_{1i})$ must satisfy for k_i and any $\tilde{N}_{1i}$:

$$u(R^*_{1i}, N^*_{1i}; k_i) - v(R^*_{1i}) \geq u(0, \tilde{N}_{1i}; 0) - v(0)$$

To simplify the analysis below, assume that when the contribution limit is equal to L, it is only optimal for those with $k_i = k_L = 0$ to contribute to the retirement account; while not contributing is optimal for those with larger fixed costs. Also, let each individual's type be denoted by their fixed cost of contributing to the retirement account (i.e. those with $k_i = k_H$ are denoted as type h individuals).

3.2 *An increase the contribution limit*

Before discussing how an (exogenous) increase in the contribution limit affects saving in this model, I first distinguish between private, public and national saving. Private saving is defined as the sum of personal and corporate saving in a given period (i.e. a year).[18] Public saving is the difference between revenues from local, state and the federal governments less expenditures. National saving is the sum of private and public saving. Given the large revenue costs associated with tax-deferred savings accounts such as IRAs and 401(k)s, a policy aimed at increasing saving is only effective

[18]Given that households own the firms (corporations) it may be unnecessary to distinguish between personal and corporate saving. However, to the extent that firms are owned by foreign individuals or firms, the private part of national saving may not be equal to the sum of personal and corporate saving.

if it increases national saving. In the discussion that follows, I focus on the effect of an increase in the contribution limit on personal saving. Total personal saving for individual i is defined as $S_i^* = R_{1i}^* + N_{1i}^*$. Following a change in the contribution limit, retirement account contributions crowd-out taxable saving if $\frac{\partial N_{1i}^*}{\partial L} < 0$ when $\frac{\partial R_{1i}^*}{\partial L} > 0$, so that $\frac{\partial S_i^*}{\partial L} < \frac{\partial R_{1i}^*}{\partial L}$.

Type l savers: Assume that some type l savers are constrained by the contribution limit L. The LCM predicts that changing the contribution limit will only affect these individuals, a prediction that is described by Proposition 1.

Proposition 1: *In the standard LCM an increase in the contribution limit: (a) increases retirement account contributions only for type l savers that are limit contributors, (b) decreases taxable account contributions only for type l savers that are limit contributors and (c) increases total personal saving only for type l savers that are limit contributors.*

Proof: See appendix. □

The mechanical increase in retirement account contributions for constrained savers is intuitive and has been modelled extensively in the literature (Venti and Wise, 1990a, 1990b; Gale and Scholz, 1994). The result that total personal saving increases for this group of savers is more subtle. Recall that a \$1 (pre-tax) contribution to the taxable account decreases period 1 consumption by \$1. However, a \$1 (pre-tax) contribution to the retirement account decreases period 1 consumption by only $\$(1-\tau_1)$. Following an increase in the limit, constrained type l savers restructure their portfolios by contributing more to the retirement account. These contributions are only partially offset by reductions in taxable account contributions, leading to higher period 1 saving. Since each dollar allocated to the retirement account only lowers period 1 consumption by $\$(1-\tau_1)$, period 1 consumption and total saving actually increases because of a lower tax liability.[19] The increase in period 1 saving also leads to higher consumption in period 2.

Type m and h savers: An increase in the contribution limit may also induce some type m and h savers to contribute to the retirement account. For those who would have saved at the initial contribution limit L in the absence of the fixed cost, an increase in the contribution limit may make it optimal to contribute an amount above the previous limit. Specifically, assume that for individuals with $k_i = k_M$ (type m savers) contributing to the retirement account is optimal when the limit is increased to $L' > L$, while not contributing is still optimal for type h savers. For these type m savers, the following two inequalities must be satisfied for any $\tilde{N_{1M}}$ and $\hat{N_{1M}}$:

[19]This increased period 1 consumption is financed by the decrease in taxes.Using the period 1 budget constraint and substituting $\frac{\partial R_{1i}^*}{\partial L} = 1$ and $\frac{\partial N_{1i}^*}{\partial L}$ for type l limit contributors gives $\frac{\partial c_{1i}^*}{\partial L} = (\tau_1 - \tau_2 + r(1-\tau_2))(u_{12} - u_{22}(1+r(1-\tau_2)) > 0$.

$$u(L, \tilde{N_{1M}}; k_M) - v(L) < u(0, N^*_{1M}; 0) - v(0)$$

and

$$u(R^{**}_{1M}, N^{**}_{1M}; k_M) - v(R^{**}_{1M}) \geq u(0, \hat{N_{1M}}; 0) - v(0)$$

where $(R^{**}_{1M}, N^{**}_{1M})$ is the the optimal level of retirement and taxable account contributions for type m savers at the new limit. Note that it must be the case that $R^{**}_{1M} \in (L, L']$. Contributing an amount equal to or below the initial contribution limit is not optimal; otherwise the individual would have saved in retirement account prior to the limit increase. Therefore, if fixed costs are the sole determinant of an increase in retirement account participation rates, it must be the case that (average) new contributions of those induced to participate are larger than the initial limit. In Section 5 I show that new contributions of those induced to participate are relatively small, suggesting that the response to the 'catch-up limit' was not limited to constrained savers.

Inspecting the participation condition for the model may offer insights into potential mechanisms that could lead to increases in retirement account participation rates without large new contributions. Abstracting from potential fixed costs, the inequality which governs retirement account participation (Case B in the appendix) is:

$$\frac{v'(R_{1i})}{u_2} \geq \tau_1 - \tau_2 + r\tau_1(1 - \tau_2)$$

This condition says that individuals do not contribute to the retirement account when the relative disutility of doing so (due to the liquidity costs of contributing) is larger than the after-tax real rate of return in the account (the net benefit associated with saving). In order for individuals to be induced to contribute to the retirement account with relatively small new contributions (i.e. flip the inequality), the relative disutility of contributing must decrease. This can happen if the disutility associated with illiquid retirement account contributions decreases or if the marginal utility of period 2 consumption increases.[20]

[20]The two-period model presented in this section precludes any dynamic responses to an increase in the contribution limit. However, simply introducing dynamics in individual savings decisions does not change two main LCM predictions. First, without fixed costs, the LCM does not predict an increase in participation rates following a limit change (see Milligan, 2003). Second, even with fixed costs, those who are induced to participate due to an increase in the contribution limit must make new retirement account contributions that are larger than the previous limit.

4 Data and Empirical Strategy

4.1 Data

Estimating how eligibility for 'catch-up limits' affects saving requires individual-level information on IRA and 401(k) contributions, taxable saving and year of birth. The SIPP is ideally suited for this analysis because it provides detailed information on retirement account ownership, contributions and asset balances, as well as relatively large sample sizes. Individuals and households in each SIPP panel are interviewed every four months (waves) for a period of three or four years. In this paper, I pool several cross-sections (waves) from the 1996, 2001 and 2004 SIPP panels.[21] Information from the 1996 panel as well as waves 1 to 4 of the 2001 panel are used to analyse the savings behavior of individuals in the years prior to the EGTRRA (1996-1998 and 2001 calendar years). Data from waves 3 to 9 of the 2001 panel and the 2004 panel are used for the post EGTRRA analysis.

Questions about IRA and 401(k) ownership and contributions are asked in waves 4, 7 and 10 (where available). The timing implies that respondents are asked about their tax-deductible contributions during the January to April period for the preceding calendar year. This coincides with tax-preparation season, the period when most contributions are made. For IRAs the questions are: (i) "Do you have an Individual Retirement Account, that is, an IRA, in your own name?", (ii) "Did you make any tax-deductible contributions to IRA accounts which applied to your 200X (199X) tax return?" and (iii) "How much were your tax-deductible contributions to IRA accounts, which applied to your 200X (199X) tax return?" Similar questions are asked for 401(k)s. I use question (ii) to examine IRA and 401(k) participation behavior while question (iii) is used to study contributions.[22]

The SIPP does not ask questions about flows in taxable assets such as bank or brokerage accounts. Therefore, as with most of the literature, non-IRA (non-401(k)) flows are computed by differencing reported year-end asset balances.[23] These imputed taxable flows provide an alternative definition of saving that differs from the responses about (deductible) contributions. In particular,

[21] The 1996 SIPP panel surveys households for a period of four years from 1996 to 1999 (12 waves) while the 2001 and 2004 panels cover a period of three years (9 waves).

[22] SIPP respondents are specifically asked about *deductible* IRA and 401(k) contributions. For individuals and households with a modified AGI above the allowable threshold to take a deduction, IRA and 401(k) contributions will be understated. In particular, these individuals may still make contributions to tax-preferred savings accounts without taking a deduction in order to benefit from the tax-free accumulation of interest that these accounts provide. Also, since each of these questions asks about deductible IRA or 401(k) contributions, savings in Roth IRAs and Roth 401(k)s are not reported even though 'catch-up limits' also apply to these accounts. For these reasons, the results may actually understate the effect of 'catch-up limits' on contribution behavior.

[23] Questions about the balances of assets and liabilities are asked in waves 3, 6, 9 and 12 (where available). For example, taxable savings for the 2002 calendar year is computed by subtracting asset balances reported in wave 6 of the 2001 SIPP panel from asset balances reported in wave 3.

Table 2: Summary Statistics: Post-EGTRRA (pooled)

		Saver Type	
Variable	All	IRA Owners	401(k) Owners
Own IRA	0.201	1.000	0.344
Made IRA Contribution	0.051	0.277	0.094
IRA Contribution	78.60	377.13	111.99
Own 401(k)	0.361	0.578	1.000
401(k) Contribution	669.34	1,505.65	1,800.26
Bank Saving	-65.06	96.74	751.63
Other Equity Saving	-45.94	-1,283.38	216.15
Stock Saving	-9,742.87	4,399.70	-7,181.31
Unsecured Debt Change	-372.09	625.98	-553.08
Age	41.71	47.35	43.88
Female	0.513	0.505	0.463
White	0.821	0.907	0.861
Black	0.117	0.041	0.084
Hispanic	0.028	0.028	0.027
HH Total Income	56,377.44	77,934.91	70,788.85
HH Earned Income	50,915.19	70,551.04	67,073.56
Personal Earned Income	22,624.21	38,392.16	41,076.50
Number of Kids < 18	0.901	0.705	0.816
High School	0.277	0.166	0.208
Some College	0.346	0.325	0.356
College Degree	0.248	0.491	0.406
N	105,339	22,240	27,206

Notes: All dollar amounts are deflated to 1996 dollars using the Bureau of Labor Statistics CPI Inflation Calculator. The sample in column 1 is (pooled) respondents from the 2001 and 2004 SIPP panels between the ages of 18 and 65. The sample in column 2 is restricted to IRA owners. The Bank Saving, Other Equity Saving, Stock Saving and Unsecured Debt Change variables for any year t are imputed by subtracting the reported assets for that variable at year t - 1 from the balance at year t. High School is a dummy variable equal to 1 if the respondent reports that they have at least a high school diploma and 0 otherwise. Some College is a dummy variable equal to 1 if the respondent reports that they began a college program but did not finish or completed a 2 year program and 0 otherwise. College Degree is a dummy variable equal to 1 if the respondent reports that their highest level of education is at least a Bachelor's Degree and 0 otherwise.

differences in asset balances between two years include any (active) contributions made by individuals as well as accumulated interest or capital gains.[24] Finally, it is also important to note that I am only able to compute taxable savings variables for the 2002, 2003 and 2005 calendar years. This is due to the timing of the SIPP waves. For the 2004 SIPP panel, for example, respondents are first asked about taxable account balances in wave 3. This corresponds to the end of the 2004 calendar year, so flows cannot be calculated for 2004. In contrast, IRA and 401(k) deductible contributions are reported for the 2002, 2004 and 2005 calendar years. To ensure that the results below are comparable with the earlier literature, I focus on five taxable savings variables in this paper: (a) bank savings, (b) stock and mutual fund (account) savings, (c) (changes in the level of) unsecured

[24] As in Gelber (2011), I also winsorize these imputed savings measures at the 5th and 95th percentiles to remove the influence of outliers. In particular, this removes observations for individuals who experience large bank/stock/mutual fund account balance changes from year to year. These large changes are likely due to major purchases, such as buying a house and do not qualitatively affect the results in Section 6.

debt held, (d) (change in the) value of vehicles owned and (e) other financial asset (OFA) savings.[25] While this list is not comprehensive, it encompasses many of the assets which may substitute for tax-preferred savings, at least in the short term. The definitions of these variables are provided in appendix A.1.

Eligibility for 'catch-up limits' in any given year is determined by an individual's year of birth. Using the respondent's date of birth recorded in the SIPP, I construct an "age" variable for each individual-year observation. Specifically, for individual i in year t, "age" is defined as $t - YOB_i$. Those whose "age" is 50 or greater after 2002 are eligible for 'catch-up limits'.

Finally, I make some sample restrictions that are common in the literature. First, I restrict the analysis to individuals between the ages of 18 and 65 in any particular year.[26] Restricting the sample to working age adults leaves 118,781 individual-year observations in the pooled post-reform sample. Next, I exclude observations from all respondents who either (a) don't answer or respond "Don't Know" to the question of whether they own an IRA, (b) responded "Refused" or "Don't Know" to the question of whether they made an IRA contribution[27] and (c) who report making a negative contribution to their IRA.[28] For the 401(k) analysis, I exclude those individuals who fit the criteria above for the relevant 401(k) questions. These restrictions lead to a sample of 105,339 individual-year observations for the post-EGTRRA period. I also deflate all savings, contribution and income variables to 1996 dollars, the first year of the pre-reform data, using the Bureau of Labor Statistics (BLS) CPI Inflation Calculator.

Column 1 of Table 2 reports sample means for the entire pooled post-EGTRRA sample, while columns 2 and 3 show means for the sub-samples of IRA and 401(k) owners. There are large differences in annual saving across individuals in these categories. Specifically, only 20% of individuals own an IRA and only 36% own a 401(k). Moreover, the average annual level of saving in these plans is quite low, suggesting that many Americans do not accumulate large amounts of financial assets for retirement. This finding is consistent with previous research (Venti and Wise, 1990a; ICI, 2012). Also, IRA owners are older, more likely to be white, have graduated from high school, have a college degree, and earn more than both the full sample and 401(k) owners. This heterogeneity in the taste and ability to save is what has made identifying the causal effect of tax-preferred savings accounts so difficult. Consistent with previous research, the likelihood of making a deductible

[25] I also explore whether tax-preferred savings crowd-out other non-IRA (non-401(k)) savings measures such as secured debt, home equity and interest earning assets not held in financial institutions, though the results are not reported for brevity. These estimates are available upon request.

[26] In sensitivity analysis to the age window selected for the main results (subsection 5.3 below) I allow the analysis to include individuals between the ages of 18 and 90.

[27] Some respondents do not answer questions (ii) and (iii), though these are almost exclusively those who report that they do not own an IRA (401(k)) and therefore cannot make tax-deductible contributions. For these cases I code the level of IRA and 401(k) contributions as $0. This does not affect the results since the analysis for the sub-sample of IRA owners (those who always respond to questions (ii) and (iii)) produces similar results, as show below.

[28] Respondents are asked a separate series of questions about withdrawals from IRA and 401(k) accounts. The results in Section 5 below are not sensitive to dropping those who report negative IRA or 401(k) contributions.

Figure 2: Distribution of IRA Contributions

(a) Pre-EGTRRA Ages 40-49

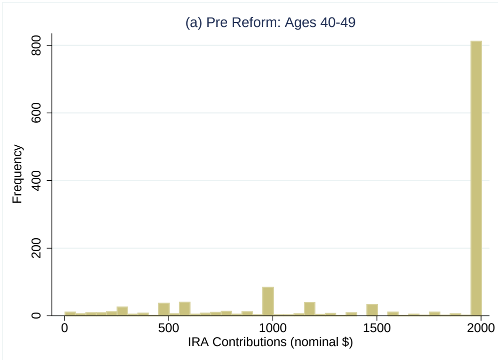

(b) Pre-EGTRRA Ages 50-59

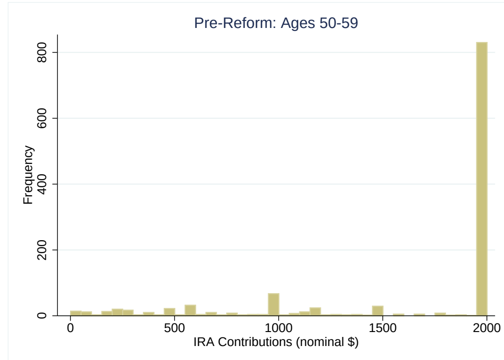

(c) Post-EGTRRA Ages 40-49

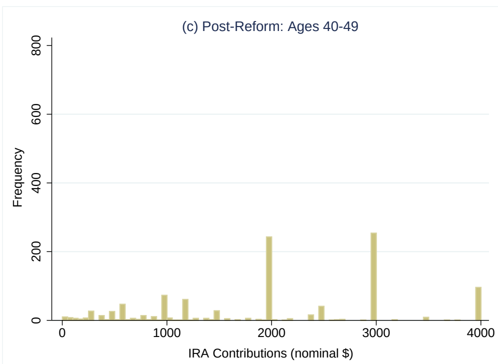

(d) Post-EGTRRA Ages 50-59

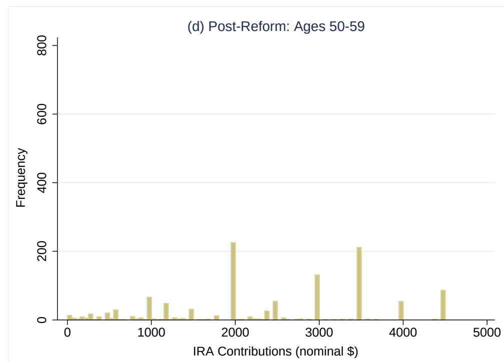

Notes: The top panel shows the distribution of reported tax-deductible IRA contributions in the pre-reform (1996-1998, 2001) period. The bottom panel shows the distribution of reported tax-deductible IRA contributions for the post-reform (2002, 2004-2005) years.

contribution is quite small among IRA owners (Holden and Bass, 2014). On average, only 28% of IRA owners (five percent of the full sample) make a deductible contribution in a given year.

Figure 2 shows the distribution of IRA contributions (for those making a positive contribution) for both 40-49 year olds and 50-59 year olds in the pre- and post-EGTRRA periods. In the pre-reform period the vast majority of IRA contributors are constrained by the $2,000 limit, as demonstrated by the large spike in the distribution. In contrast, in the post-reform years, fewer individuals are constrained by the contribution limit. Interestingly, after 2002 there is still a large spike in the distribution at $2,000 for both those over and under the age of 50. This observation is consistent with the inertia in retirement savings behavior that is well documented in the literature (Madrian and Shea, 2001; Choi et al., 2004; Chetty et al., 2014).

4.2 Empirical Strategy

4.2.1 How do higher contribution limits affect contributions to IRAs and 401(k)s?

One of the challenges in identifying the causal savings effect of tax-advantaged accounts is finding exogenous variation in the eligibility of individuals to participate in these plans. Before describing the identification strategy used in this paper, it is useful to describe an "ideal experiment" to analyze the savings effect of an increase in the contribution limit. Such an ideal experiment could be achieved by either randomizing the contribution limit for a sample of individuals or by increasing the contribution limit for some, leaving the initial limit unchanged for others. Differences in contributions could then be interpreted as the causal effect of changing the limit.

In the United States, however, most changes to IRA and 401(k) eligibility has either been universal or correlated with potential determinants of individual tastes for saving, such as workplace pension status or AGI. In this paper I use quasi-experimental variation in contribution limits induced by the 'catch-up limit' provision to identify the causal effect of a limit change on IRA and 401(k) saving. As discussed in Section 2, 'catch-up limits' allow individuals over the age of 50 to make larger IRA and 401(k) contributions than those under the age of 50. This provides an exogenous source of variation in the ability of individuals to contribute to tax-deferred savings accounts, resembling a randomized experiment for individuals in a neighborhood of age 50.

I estimate the effect of eligibility for 'catch-up limits' on IRA and 401(k) contributions using a (sharp) regression discontinuity design (RDD). Identification of the savings effect of this policy only requires that the counterfactual IRA (or 401(k)) savings function (the savings function in the absence of a *differential change* in contribution limits) is continuous at age 50. Although this assumption is untestable, I present evidence below that supports its plausibility.

The main estimating equation for the RDD is:[29]

$$R_{it} = \alpha + \beta over50_{it} + h(\tilde{age}_{it}) + Z_{it}^{'}\gamma + \delta_t + u_{it} \tag{2}$$

where R_{it} denotes the deductible IRA (or 401(k)) contributions of individual i in year t, $over50_{it}$ is a dummy variable equal to one if the individual is eligible for 'catch-up limits', $\tilde{age}_{it}$ is the normalized assignment variable for individual i, $h(\tilde{age}_{it})$ is a flexible polynomial in (normalized) age that varies on each side of the cutoff, Z_{it} is a vector of covariates for individual i, δ_t are year fixed effects and u_{it} is the residual.[30] The inclusion of covariates and year fixed effects in equation (2) is not required for identification but appear in some specifications to improve precision. The baseline

[29]Equation (2) can be thought of as a reduced form equation representing the solution for optimal retirement account savings for the model in Section 3, where saving is a smooth function of age.

[30]Recall that age is defined as the age of an individual at the end of the calendar year. As is standard in RDD, this variable is normalized so that 0 denotes individuals that are 50 years old.

covariates are race dummies (white, black and hispanic (the excluded category is Asian/other)), demographic dummies (female and married), number of children under the age of 18, (personal) earnings (deflated to 1996 dollars), as well as a set of education dummies. In specifications without covariates, the parameter α represents the estimate for the counterfactual level of contributions at age 50 (those just below the cutoff). The parameter β captures the casual effect of being eligible for 'catch-up limits' on contributions or participation rates. All regressions weight observations by the inverse sampling probabilities (person weights) provided in the SIPP.[31] Following Card and Lee (2008), standard errors are clustered at the age level to account for any uncertainty in the estimated effect of the 'catch-up limit' provision due to misspecification of the age polynomial.

The parameter β is a local intent-to-treat (ITT) estimate of the effect of *eligibility* for 'catch-up limits' on retirement savings. Regression discontinuity designs only identify causal effects in a neighborhood of the assignment cutoff. Therefore, β captures the causal effect of eligibility for a limit change for individuals at a prime savings age. This estimate should be thought of as an ITT for several reasons. Since individuals with incomes above the deductibility thresholds can make non-deductible contributions (up to the limit) and because the reported IRA contributions of these individuals is $0 in the SIPP, β may understate the effect of eligibility for 'catch-up limits'. Also, some over the age of 50 may have been unaware that they were able to make larger contributions in the years immediately following the reform.

β also incorporates any potential dynamic responses by individuals both over and under 50. For example, those under the age of 50 could delay planned contributions until they become eligible to make larger contributions, as in Milligan (2003). However, most pre-EGTRRA IRA contributors were constrained by the limit (see Figure 2), suggesting that the group of savers who might delay contributions because they were both unconstrained before the EGTRRA and expected to be constrained in the future is likely small. Moreover, the "use-it-or-lose-it" mechanism described in Milligan (2003) does not predict an increase in participation rates, and delaying planned IRA or 401(k) contributions would mean forging the benefit from the tax-free accumulation of interest.

Alternatively, savers who planned to make retirement account contributions in their late 50s may have been induced to bring these contributions forward in response to eligibility for 'catch-up limits'. As a result, observed increases in IRA or 401(k) contributions immediately following the reform may have crowded-out future contributions for these savers. Though beyond the scope of this study, any potential dynamic responses imply that increases in contributions (i.e. a positive estimate for β) in the short term may not represent long term increases in saving. However, recent research finds that many individuals fail to respond to even contemporaneous changes in incentives to save in retirement accounts (Chetty et al., 2014), suggesting that the potential dynamic responses due to re-optimizing saving over the life-cycle may be small.

[31] OLS estimates (without weights) are reported in appendix Table A.4 as a comparison. Overall, both the weighted and unweighted estimates are qualitatively and quantitatively similar.

In RDD it is standard to investigate whether there is manipulation of the assignment variable. In this application, such manipulation is unlikely since eligibility for 'catch-up limits' is a deterministic function of an individual's year of birth. Since deductible IRA and 401(k) contributions are claimed when filing income taxes, individuals would have to lie about their age to the IRS about in order to manipulate the assignment variable. Nonetheless, Figure A.1 in the appendix investigates this possibility by plotting a histogram of the age distribution in the post-EGTRRA sample around the eligibility cutoff. There is no evidence of a spike in the number of individuals in the sample just above the cutoff.[32] Figure A.2 (appendix) tests for discontinuities in other baseline covariates.[33] Reassuringly, almost all covariates are smooth around the age 50 threshold. The one exception is the college completion dummy variable. Individuals just over the age of 50 are approximately three percentage points more likely to have a college degree than those just under the age of 50. If college completion is positively correlated with tastes for saving (perhaps due to increased financial literacy or smaller discount rates), then estimates of β may be biased upward. This increase in the likelihood of having a college degree for individuals in this age group is probably due to the increased likelihood of enrolling in and remaining in college to avoid military service during the Vietnam War period. In the main results below, however, I show that the inclusion of this covariate does not impact the estimates of β either qualitatively or quantitatively.

4.2.2 Do IRA and 401(k) contributions crowd-out taxable saving?

I estimate the extent to which increases in IRA and 401(k) contributions are achieved by reducing taxable saving by estimating two-stage least squares (2SLS) regressions of the following form:

$$N_{it} = \mu_0 + \mu_1 R_{it} + \mu_2 h(\tilde{age}_{it}) + Z_{it}^{'}\mu_3 + \delta_t + \epsilon_{it} \tag{3}$$

where N_{it} denotes the measure of taxable saving for individual i in year t and ϵ_{it} is the residual. The remaining variables have the same definitions as in equation (2). Tax-preferred saving, R_{it}, is instrumented with the $over50_{it}$ dummy variable (i.e. eligibility for 'catch-up limits'). In order for the instrument to be valid, eligibility for 'catch-up limits' must be correlated with IRA contributions and uncorrelated with the residual in equation (3). If the identification assumption for the RDD holds, the exclusion restriction is also likely to be satisfied. In order to be consistent with the

[32] I formally test for a discontinuity in the assignment variable by collapsing the data into age bins and regressing the number of individuals (in the pooled post-EGTRRA sample) in each age bin on a polynomial in (normalized) age and a dummy variable for those over the age of 50. Each age bin observation is weighted by the relative number of individuals in that age bin. With a bandwidth of 20 and a quadratic spline in (normalized) age I find no evidence of a discontinuity at age 50. I obtain similar results using an age window of 10 on either side of the cutoff and a linear spline, the specification used in the main results below.

[33] Table A.1 in the appendix formally tests this be estimating linear splines for the various covariates for individuals within 10 years of the cutoff. Results of this confirm the visual evidence in Figure A.2.

previous literature, I use the imputed IRA savings variable (i.e. IRA balance at the end of year t minus account balance at year $t-1$) for R_{it} in the analysis in Section 6. I also estimate equation (3) using the reported tax-deductible contributions for R_{it} in Table A.3 of the appendix.

The parameter μ_1 captures the crowd-out parameter of interest; on average a $1 increase in IRA saving leads to a μ_1 increase in taxable saving. A negative estimate of μ_1 would imply that increases in IRA contributions are (at least partially) achieved by reshuffling assets from taxable accounts to IRAs. In the extreme case of $\mu_1 < -1$, increases in IRA contributions lead to decreases in personal saving. Recall that the stylized model in Section 3 predicts that μ_1 is negative but larger than -1 for limit contributors.

5 The effect of eligibility for 'catch-up limits' on retirement saving

This section presents estimates for the effect of raising contribution limits on saving in retirement accounts. In subsections 5.1 and 5.2, I report and discuss the effect of eligibility for 'catch-up limits' on IRA and 401(k) contributions respectively. The section concludes with an analysis of the robustness of the main results.

5.1 IRAs

Figure 3 illustrates the regression discontinuity design by plotting IRA contributions against age. Each circle represents the mean IRA contribution for a given age bin, with age 50 normalized to zero[34]. A linear spline is estimated on each side of the cutoff for individuals between the ages of 40 and 59. The solid lines represent the fitted values from these regressions and the dashed lines the corresponding 95% confidence intervals. Figure 3a plots (unconditional) average IRA contributions against age during the pre-EGTRRA years (1996-1998, 2001), while Figure 3b plots the same relationship for the post-EGTRRA years (2002, 2004-2005).

The identification assumption for the RDD requires that tastes for saving be continuous in age at 50 in the absence of 'catch-up limits'. While there is no theoretical reason why saving (tax-preferred or taxable) would increase discontinuously at age 50, if this age is "special" in any way the identification assumption would be violated. One way to test the plausibility of this assumption is to check whether IRA contributions "jump" at age 50 before 'catch-up limits' are introduced. Figure 3a illustrates this placebo test; there is no evidence of a discontinuity at age 50.[35] However,

[34] Recall that the assignment variable, age_{it}, for individual i in year t is defined as $t - YOB_i$ where YOB_i is year of birth.

[35] This relationship is estimated formally with the regression described by equation (2). The RDD estimate is -$8 and is not statistically significant. The regression results for all other placebo tests discussed below are available upon request.

Figure 3: Mean IRA Contributions by Age

(a) Pre-EGTRRA (b) Post-EGTRRA

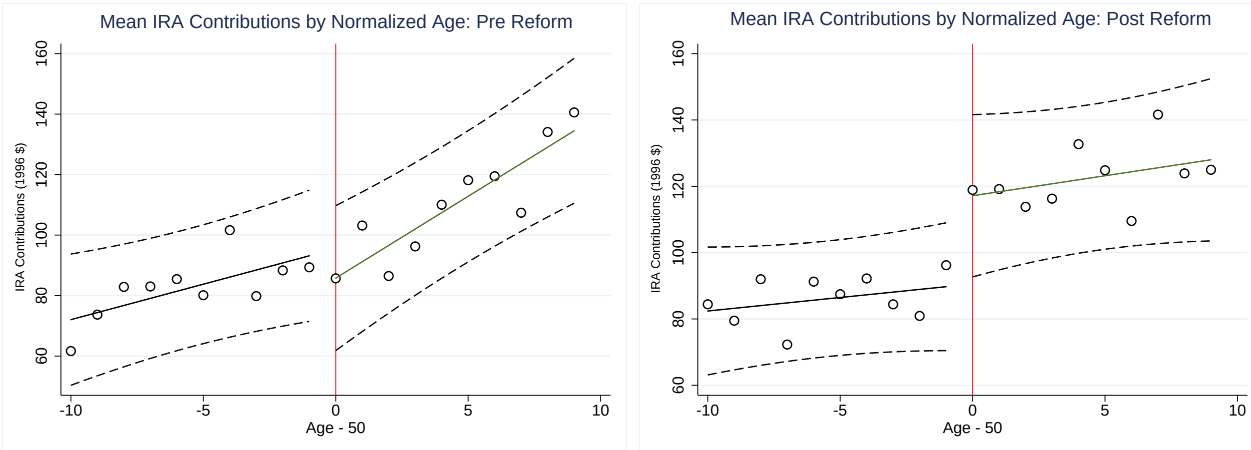

Notes: All dollar amounts are deflated to 1996 dollars using the Bureau of Labor Statistics CPI Inflation Calculator. The sample is restricted to all individuals between the ages of 40 and 59 during the pre and post-EGTRRA years. Figure 3a plots (unconditional mean) IRA contributions against the normalized assignment variable (age) for the 1996-1998 and 2001 pre-EGTRRA years. Figure 3b plots (unconditional mean) IRA contributions against the normalized assignment variable (age) for the 2002 and 2004-2005 post-EGTRRA years. Linear spline regressions are estimated on each side of the cutoff and the corresponding 95 percent confidence intervals are represented by the dashed lines.

in the post-EGTRRA years (Figure 3b) average IRA contributions jump at age 50 from approximately \$90 to about \$120.

This relationship is estimated formally by equation (2) and the results are reported in panel A of Table 3. A linear spline is estimated on each side of the cutoff, allowing the polynomial in (normalized) age to differ for those over 50. All regressions weight individuals by their inverse sampling probabilities in the SIPP. I restrict the analysis to individuals between the ages of 40 and 59 (bandwidth of 10). Selecting this bandwidth has several advantages. First, IRAs feature early withdrawal penalties for those under the age of 59.5. As a result, the contribution behavior of those 60 years of age and older may be very different than that of younger savers. Second, a bandwidth of 10 is more conservative than the optimal bandwidth of 26 selected by the Imbens and Kalyanaraman (2012) (IK) procedure. Interestingly, in the regressions without sample weights reported in the appendix, the IK optimal bandwidth is 9, suggesting that an age window of 10 is reasonable.[36] Moreover, in subsection 5.3 I show that the main results do not vary with the bandwidth choice.

The top row in columns 1-3 report the estimates for α, the counterfactual level of IRA con-

[36] The SIPP under-samples high-income individuals and households. Since the variance of IRA contributions rises with income, the regressions with sample weights which place relatively more weight on higher income individuals will have a larger (residual) variance. The IK procedure which selects the bandwidth that minimizes the MSE will therefore lead to a larger bandwidth. For the unweighted regressions that restrict the analysis to IRA owners, the optimal bandwidth is 11.

tributions at age 50. The second row reports the RDD estimates for the effect of eligibility for ‘catch-up limits’. The specification in column 1 does not include year fixed effects or covariates. In columns 2 and 3, year fixed effects and covariates are added. Eligibility for ‘catch-up limits’ is estimated to increase (unconditional) average IRA contributions by between $19 and $23 (about 25%) at age 50. Estimates with or without year fixed effects and covariates are similar in magnitude and statistically significant at conventional levels.

Table 3: Effect of ‘Catch-up Limits’ on IRA Contributions and Participation Rates

	(1)	(2)	(3)	(4)	(5)	(6)
		Contributions			Participation Rates	
A. Full Sample = 51,725						
Control Mean	93.99	79.15	-55.00	0.045	0.040	-0.012
	(6.89)***	(8.18)***	(13.78)***	(0.004)***	(0.005)***	(0.006)***
$Over50$	22.58	22.82	18.79	0.011	0.011	0.009
	(7.30)***	(7.27)***	(6.95)***	(0.004)***	(0.004)***	(0.004)**
Year FE	N	Y	Y	N	Y	Y
Covariates	N	N	Y	N	Y	Y
B. IRA Owners = 12,261						
Control Mean	398.70	386.97	273.32	0.190	0.193	0.252
	(22.96)***	(33.57)***	(102.45)***	(0.016)***	(0.019)***	(0.045)***
$Over50$	77.60	78.08	72.66	0.036	0.036	0.033
	(22.58)***	(28.99)**	(27.83)**	(0.017)**	(0.017)*	(0.017)*
Year FE	N	Y	Y	N	Y	Y
Covariates	N	N	Y	N	N	Y
C. Positive Contributors = 2,501						
Control Mean	2,098.50	2,018.66	1,123.83	1.000	1.000	1.000
	(48.35)***	(61.20)***	(268.03)***	(0.000)***	(0.000)***	(0.000)***
$Over50$	1.04	6.68	-37.34	0.000	0.000	0.000
	(75.47)	(75.33)	(74.84)	(0.000)	(0.000)	(0.000)
Year FE	N	Y	Y	N	Y	Y
Covariates	N	N	Y	N	N	Y

Notes: All dollar amounts are deflated to 1996 dollars using the Bureau of Labor Statistics CPI Inflation Calculator. The dependent variable in columns 1-3 is a respondent’s reported IRA contribution for the 2002, 2004 or 2005 calendar year. The dependent variable in columns 4-6 is a dummy variable equal to 1 if the respondent reported making a deductible IRA contribution for the 2002, 2004 or 2005 calendar year. $Over50$ is a dummy variable equal to 1 if the respondent is eligible for ‘catch-up limits’. In all columns, the sample is restricted to individuals between the ages of 40 and 59 (age window of 10) and regression results are for a polynomial of degree one in age that is allowed to vary on either side of the cutoff. The sample in panel A is all respondents between the ages of 40 and 59. The sample in panel B is all IRA owners. The sample in panel C is all individuals who contribute a positive amount to their IRA. Each observation is weighted by its inverse sampling probability in the SIPP. The covariates (as defined in Table 2) are: White, Black, Hispanic, Female, Married, Personal Earned Income, Number of Kids Under the age of 18, High School, Some College and College Degree. Standard errors are clustered at the age level and appear in brackets. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure 4 plots the likelihood of making a positive IRA contribution against age around the ‘catch-up limit’ eligibility cutoff. The dependent is equal to one if the respondent reports making a deductible IRA contribution in the relevant year and zero otherwise. As is the case with contributions, I cannot reject the null hypothesis that the IRA participation rate is continuous at age 50 before ‘catch-up limits’ are introduced (statistically insignificant point estimate of 0.000

Figure 4: Mean IRA Participation Rates by Age

(a) Pre-EGTRRA

(b) Post-EGTRRA

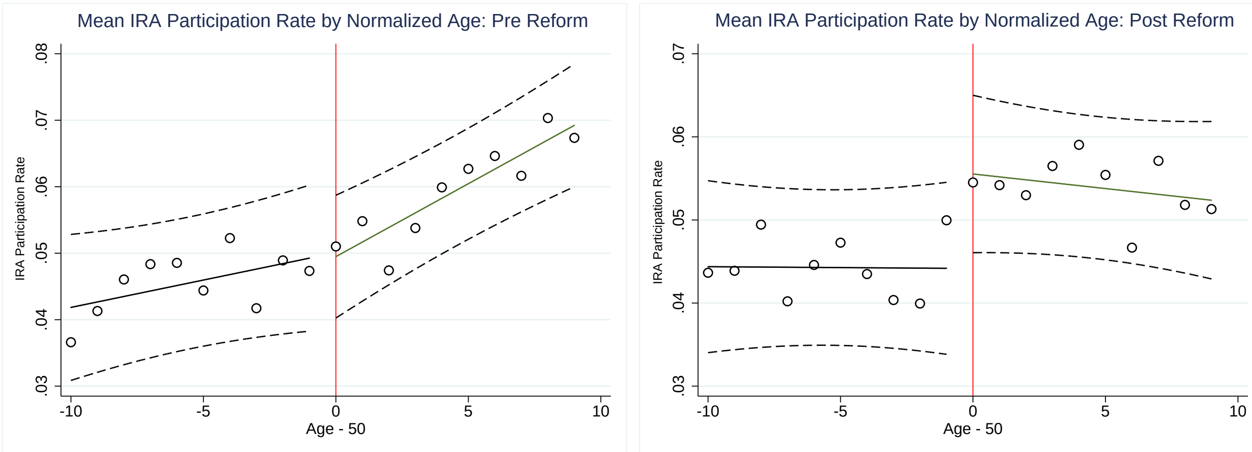

Notes: The sample is restricted to all individuals between the ages of 40 and 59 during the pre and post-EGTRRA years. Figure 4a plots (unconditional mean) IRA participation rates against the normalized assignment variable (age) for the 1996-1998 and 2001 pre-EGTRRA years. Figure 4b plots (unconditional mean) IRA participation rates against the normalized assignment variable (age) for the 2002 and 2004-2005 post-reform years. Linear spline regressions are estimated on each side of the cutoff and the corresponding 95 percent confidence intervals are represented by the dashed lines.

for β). However, during the post-EGTRRA years, the likelihood of making a positive contribution increases by about one percentage point from approximately 4.5 to 5.5 percent at age 50. In columns 4-6 of panel A, this relationship is estimated using regressions similar to those in columns 1-3. The estimate for β is 1.1 percentage points (24%) without year fixed effects or covariates and is statistically significant at the 5 percent level. The coefficient estimate is unchanged when year fixed effects (column 5) and baseline covariates (column 6).

Restricting the analysis to the sub-sample of IRA owners[37] yields qualitatively similar results. Columns 1-3 of panel B report estimates for the effect of eligibility for 'catch-up limits' on contributions for IRA owners.[38] The average IRA contribution for those just below the cutoff is \$399. Eligibility for 'catch-up limits' causes contributions to jump by a statistically significant \$78 (20%). The estimate for β is similar in magnitude and remains statistically significant when both year fixed effects and covariates are added. In columns 4-6 of panel B I explore whether eligibility for 'catch-up limits' causes IRA participation rates to increase for account owners. The coefficient estimate for α is 0.19, suggesting that about one in five IRA owners just below age 50 makes a deductible contribution in a given year. Eligibility for 'catch-up limits' increases the likelihood of making a positive contribution by a statistically significant 3.6 percentage points (19%).

[37] Specifically, I restrict the analysis to respondents who answer "yes" to the question "Do you have an Individual Retirement Account, that is, an IRA, in your own name?"

[38] In unreported results, I cannot reject the null hypothesis that deductible IRA contributions are continuous at age 50 for IRA owners during the pre-EGTRRA years (statistically insignificant point estimate for β of -\$24).

As discussed earlier, in the standard LCM we do not expect that a change in the contribution limit will affect agents for whom the limit isn't binding. However, fixed costs of contributing may cause some agents not to contribute if the benefit from doing so is too low. Expanding IRA eligibility by increasing the contribution limit may make contributing optimal for these agents. If fixed costs are driving the observed increase in IRA participation rates, then new contributions by these agents must be greater than the initial limit of \$2,000; otherwise contributing prior to the limit change would have been optimal.

I test whether the increased likelihood of making an IRA contribution can be explained by fixed costs in two ways. The first test exploits the fact that the jump in IRA contributions at age 50 (the ITT estimate for β) is a weighted average of (a) the increased likelihood of making a contribution and (b) the mechanical increase in contributions from those constrained by the pre-EGTRRA limit. By decomposing the ITT estimate into these two responses and using the RDD coefficient estimates from Table 3, I recover an estimate for the average new contributions of those induced to participate. To simplify notation, define the 'catch-up limit' eligibility indicator $D_{it} = 1$ if $over50_{it} = 1$ and $D_{it} = 0$ if $over50_{it} = 0$. Using formula (3.4.4) in Angrist and Pischke (2009, p. 98), the ITT $\beta^C = E[R_{it}|D_{it} = 1, age_{it} = 50] - E[R_{it}|D_{it} = 0, age_{it} = 50]$ can be written as

$$\beta^C = \underbrace{\left(P[R_{it} > 0|D_{it} = 1] - P[R_{it} > 0|D_{it} = 0]\right)}_{\text{Participation Effect}} E[R_{it}|R_{it} > 0, D_{it} = 1]$$

$$+ \underbrace{\left(E[R_{it}|R_{it} > 0, D_{it} = 1] - E[R_{it}|R_{it} > 0, D_{it} = 0]\right)}_{\textit{Conditional-on-positive (COP) Effect}} P[R_{it} > 0|D_{it} = 0]$$

For convenience, I suppress the fact that each of the expectations on the right side of this expression are also conditional on $age_{it} = 50$. Thus, β^C is a weighted average of the increase in IRA participation rates (the participation effect) and the increase in contributions at age 50 by those making a positive contribution (the COP effect). If the RDD is valid, then the difference $P[R_{it} > 0|D_{it} = 1] - P[R_{it} > 0|D_{it} = 0] = \beta^P$ is a consistent estimate of the participation effect, and the probability that those just under 50 make an IRA contribution is $P[R_{it} > 0|D_{it} = 0] = \alpha$. However, it is not possible to estimate the average new contributions of those induced to participate directly from this expression. This is because the average contribution by those at age 50 who contribute a positive amount, $E[R_{it}|R_{it} > 0, D_{it} = 1]$, is itself a weighted average of the new contributions by those induced to participate and the contributions of those constrained by the limit.

Let $s = M$ denote the (marginal) savers induced to contribute to an IRA by the introduction of 'catch-up limits' and $s = I$ denote the inframarginal savers who would have contributed in the absence of the reform. The fraction of IRA contributors that are marginal savers is $\frac{\beta^P}{\alpha+\beta^P}$. Thus, the average IRA contribution by those just over 50 who save a positive amount, $E[R_{it}|R_{it} > 0, D_{it} = 1]$, is $\psi E[R_{it}|R_{it} > 0, D_{it} = 1, s = M] + (1-\psi)E[R_{it}|R_{it} > 0, D_{it} = 1, s = I]$, where the expectation $E[R_{it}|R_{it} > 0, D_{it} = 1, s = I]$ is the average contribution by inframarginal savers. The moment of interest, $E[R_{it}|R_{it} > 0, D_{it} = 1, s = M]$, is the average new contribution by those induced to participate. Substituting this into the expression above and simplifying gives[39]

$$\beta^C = \beta^P E[R_{it}|R_{it} > 0, D_{it} = 1, s = M] + \alpha\Big(E[R_{it}|R_{it} > 0, D_{it} = 1, s = I] - E[R_{it}|R_{it} > 0, D_{it} = 0]\Big)$$

Although the difference $E[R_{it}|R_{it} > 0, D_{it} = 1, s = I] - E[R_{it}|R_{it} > 0, D_{it} = 0]$ cannot be estimated consistently since $E[R_{it}|R_{it} > 0, D_{it} = 1, s = I]$ is unobserved, it is bounded below by \$0. It is also bounded above by \$500 since 'catch-up limits' allow those over 50 to contribute at most \$500 more than those under 50 for the 2002-2005 period.[40] Assuming that 'catch-up limits' do not affect the saving decisions of those currently ineligible implies that $E[R_{it}|R_{it} > 0, D_{it} = 0]$ can be replaced by the estimate for the average IRA contribution by those just under age 50. As described earlier, the parameters β^C, β^P and α can be estimated consistently if the RDD is valid.

I substitute the RDD estimates for these parameters from Table 3 to recover an estimate of $E[R_{it}|R_{it} > 0, D_{it} = 1, s = M]$ for various values of the increase in IRA contributions for inframarginal savers, $E[R_{it}|R_{it} > 0, D_{it} = 1, s = I] - E[R_{it}|R_{it} > 0, D_{it} = 0]$. Standard errors are calculated using the delta method. This procedure suggests that average new contributions are relatively small and much lower than the initial IRA contribution limit of \$2,000. When contributions for inframarginal savers increase by \$250, the estimate for $E[R_{it}|R_{it} > 0, D_{it} = 1, s = M]$ implies that new contributions by those induced to participate is \$1,057; the 95% confidence interval rules out contributions larger than \$1,867. In the extreme case when contributions do not increase for inframarginal savers (i.e. $E[R_{it}|R_{it} > 0, D_{it} = 1, s = I] - E[R_{it}|R_{it} > 0, D_{it} = 0]$ is equal to \$0), the estimate for $E[R_{it}|R_{it} > 0, D_{it} = 1, s = M]$ is only \$2,120. However, this is unlikely since two thirds of IRA contributors were constrained by the pre-EGTRRA limit, suggesting a mechanical savings increase for at least some inframarginal individuals.

To complement this analysis, I also use information in the SIPP about how long individuals have contributed to an IRA to assess whether those induced to participate make large new deductible contributions. Specifically, I calculate the average IRA contributions for those between the ages of 50 and 52 who report that they have contributed to an IRA for less than one year. This procedure is not without limitations however. Notably, the average IRA contribution by those just over age 50 who have contributed for less than a year includes by those induced to participate by 'catch-up limits' as well as those who would have contributed anyway. With this in mind, the average contribution for these respondents is \$1,694. Despite the relatively small sample size for

[39]This formula is derived in Appendix A.3.

[40]This implicitly assumes that (i) counterfactual (in the absence of 'catch-up limits') IRA contributions are continuous at age 50 and (ii) that eligibility for 'catch-up limits' does not lower IRA contributions, conditional on participation.

this group (66 respondents across all three post-EGTRRA years), varying the age window produces similar estimates. Overall, there is no strong evidence indicating that average new contributions are greater than $2,000 for those induced to participate. This conclusion is consistent with the negligible estimates for the "jump" in IRA contributions at age 50 for the sub-sample of positive contributors. In panel C of Table 3, I estimate equation (2) on this sub-sample and find no evidence of a discontinuity in tax-deductible contributions. While the estimates in panel C cannot be interpreted as causal effects,[41] they are consistent with negative selection arising from those with relatively lower tastes for saving being induced to make a positive IRA contribution.

There are heterogeneous responses to eligibility for 'catch-up limits'. Estimated treatment effects are largest for individuals with both high incomes and low levels of non-IRA and non-401(k) assets. In particular, the likelihood of making a positive contribution jumps by 5.9 percentage points (36%) and average contributions jump by $137 (45%) at age 50 for IRA owners with incomes above the median and who do not own any stocks or mutual funds outside of their IRA. Estimated effects for individuals with lower than median incomes and higher levels of taxable assets are much smaller in magnitude, and in some cases are statistically indistinguishable from zero. It is plausible that financial institutions looking to increase assets under management would direct advertising towards those with weak savings histories and relatively high incomes.

Increases in average contributions of about 20% are smaller in magnitude than the 45% simulated response to an IRA limit increase in Venti and Wise (1990a, 1990b).[42] As discussed earlier, the response to a change in the contribution limit in these papers is entirely driven by a mechanical increase in saving among those constrained by the limit. As a result, the results are not directly comparable.

5.2 401(k) Accounts

Table 4 reports estimates for the effect of 'catch-up limit' eligibility on 401(k) contributions.[43] I restrict the analysis to individuals within 10 years of age of the cutoff (age window of 10) to ensure that the results are comparable with the IRA analysis in the previous subsection.[44] 401(k) contri-

[41]This is due to the fact that an increase in IRA participation rates may mean that the group of savers over the age of 50 is not comparable with those under the age of 50 even if the RDD is valid.

[42]The estimated effects documented in this section are also much smaller than those found in experimental studies that ask how contributions to tax-preferred savings accounts vary with price incentives, notably the effective match rates in Duflo et al. (2006) and Saez (2009).

[43]Figure A.5 in the appendix visually displays this result.

[44]The sample size for the 401(k) estimates reported below is much smaller than that for IRAs. This is because many more respondents don't answer or respond "Don't Know" to the question of whether they own a 401(k) than the same question for IRAs (105,339 individual-year observations for IRA ownership vs. 71,002 observations for 401(k) ownership in the full age 18-65 sample) Unless the rate of non-respondence varies for those just below relative to those just above the age 50 cutoff, this should not bias the estimates reported below. I test whether the rate of non-respondence varies discontinuously around age 50 and find no evidence of a jump in the probability of non-

butions are estimated to increase by \$80 (8%) at age 50 during the pre-EGTRRA years, though this estimate is only significant at the 10 percent level. In panel A column 1, the average 401(k) contribution for individuals just below the age of 50 is \$1,139. Eligibility for 'catch-up limits' is estimated to increase 401(k) contributions by a statistically insignificant \$5 (standard error of \$57). Similarly, when year fixed effects and covariates are added in columns 2 and 3, the coefficient estimates remain small in magnitude and are always statistically insignificant.

Columns 4-6 report estimates for the effect eligibility for 'catch-up limits' on the likelihood of making a 401(k) contribution. The unconditional participation rate for those just below the age of 50 is 0.32. The estimate for β is -0.021 (-6%) and is statistically indistinguishable from zero. The point estimate remains similar in magnitude when year fixed effects are added in column 5. However, when covariates are added in column 6, the point estimate decreases slightly to -0.026 (-8%) and is statistically significant at the 5 percent level. This suggests that eligibility for 'catch-up limits' may have slightly decreased the likelihood of contributing to a 401(k), though by much less than the increase in the likelihood of making an IRA contribution.

Panels B and C explore this further by restricting the analysis to 401(k) owners and positive contributors respectively. Among 401(k) owners just below the age of 50, the average contribution is \$3,538. Contributions are estimated to "jump" by \$239 (7%), though this estimate is highly sensitive to the inclusion of covariates. In column 3, the estimate for β decreases to \$85 and is statistically indistinguishable from zero. Among 401(k) owners, eligibility for 'catch-up limits' has a negligible effect on the likelihood of making a contribution. This is not surprising since almost all 401(k) owners make deductible contributions, typically through regular payroll deductions.

The lack of a response for 401(k) contributions and participation rates is likely due to two factors. First, the number of savers constrained by the pre-EGTRRA limit of \$10,500 is relatively small, making it difficult to detect a large mechanical effect on contributions. Specifically, only 10 percent of pre-EGTRRA 401(k) contributors were constrained by the \$10,500 limit, while approximately two thirds of IRA contributors were constrained. Furthermore, if the increase in observed IRA participation rates at age 50 can be (at least partially) attributed to successful targeted advertising by financial institutions, one would expect a smaller estimated effect for 401(k)s unless employers also attempted to raise contribution rates. Overall, the data suggest that eligibility for 'catch-up limits' leads to insignificant effects on 401(k) contributions. Therefore, the robustness checks and discussion below will focus only on IRAs.

respondence (result available upon request). Also, the empirical findings in this subsection are robust to recoding missing 401(k) contribution observations as \$0 for those who responded to the IRA ownership question, but not to the 401(k) ownership question.

Table 4: Effect of 'Catch-up Limits' on 401(k) Contributions and Participation Rates

	(1)	(2)	(3)	(4)	(5)	(6)
		Contributions			Participation Rates	
A. Full Sample = 29,365						
Control Mean	1,139.02	780.18	-518.52	0.322	0.242	-0.018
	(38.40)***	(47.49)***	(83.24)***	(0.010)***	(0.011)***	(0.020)
$Over50$	5.29	11.92	-41.93	-0.020	-0.018	-0.026
	(57.41)	(56.78)	(47.26)	(0.013)	(0.013)	(0.010)**
Year FE	N	Y	Y	N	Y	Y
Covariates	N	N	Y	N	Y	Y
B. 401(k) Owners = 9,041						
Control Mean	3,537.64	3,226.71	537.61	0.998	0.999	1.000
	(55.03)***	(79.89)***	(265.90)*	(0.001)***	(0.001)***	(0.001)***
$Over50$	238.88	234.74	84.68	0.001	0.001	0.001
	(77.83)***	(79.71)**	(81.88)	(0.001)	(0.001)	(0.001)
Year FE	N	Y	Y	N	Y	Y
Covariates	N	N	Y	N	N	Y
C. Positive Contributors = 9,037						
Control Mean	3,543.81	3,231.30	538.88	1.000	1.000	1.000
	(57.90)***	(81.89)***	(265.80)*	(0.000)***	(0.000)***	(0.000)***
$Over50$	234.08	229.82	78.22	0.000	0.000	0.000
	(79.60)***	(81.66)**	(84.20)	(0.000)	(0.000)	(0.000)
Year FE	N	Y	Y	N	Y	Y
Covariates	N	N	Y	N	N	Y

Notes: All dollar amounts are deflated to 1996 dollars using the Bureau of Labor Statistics CPI Inflation Calculator. The dependent variable in columns 1-3 is a respondent's reported 401(k) contribution for the 2002, 2004 or 2005 calendar year. The dependent variable in columns 4-6 is a dummy variable equal to 1 if the respondent reported making a deductible 401(k) contribution for the 2002, 2004 or 2005 calendar year. $Over50$ is a dummy variable equal to 1 if the respondent is eligible for 'catch-up limits'. In all columns, the sample is restricted to individuals between the ages of 40 and 59 (age window is 10) and regression results are for a polynomial of degree one in age that is allowed to vary on either side of the cutoff. The sample in panel A is all respondents between the ages of 40 and 59. The sample in panel B is all 401(k) owners. The sample in panel C is all individuals who contribute a positive amount to their 401(k). Each observation is weighted by its inverse sampling probability in the SIPP. The covariates (as defined in Table 2) are: White, Black, Hispanic, Female, Married, Personal Earned Income, Number of Kids Under the age of 18, High School, Some College and College Degree. Standard errors are clustered at the age level and appear in brackets. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

5.3 Robustness Checks

5.3.1 Bandwidth, Functional Form and Placebo Cutoffs

Examining the sensitivity of estimates to the selection of bandwidth and functional form choice is standard in RDD (Lee and Lemieux, 2010). In Figure 5, I report RDD estimates for the effect of eligibility for 'catch-up limits' on contributions and participation rates for IRA owners when the age window on each side of the cutoff varies from 4 (46 to 53 year olds) to 20 (30 to 69 year olds). Figure 5a reports the RDD estimates for IRA contributions and participation rates are displayed in Figure 5b. The solid red line in each figure indicates the bandwidth of 10 used in the main results. Each dot corresponds to the coefficient estimate for a particular bandwidth choice and the blue lines denote the 95% confidence intervals. All estimates are from linear splines with year fixed effects. In Figure 5a, coefficient estimates for IRA contributions are very similar in magnitude to

Figure 5: Bandwidth Sensitivity (IRA Owners)

(a) Contributions　　(b) Participation Rates

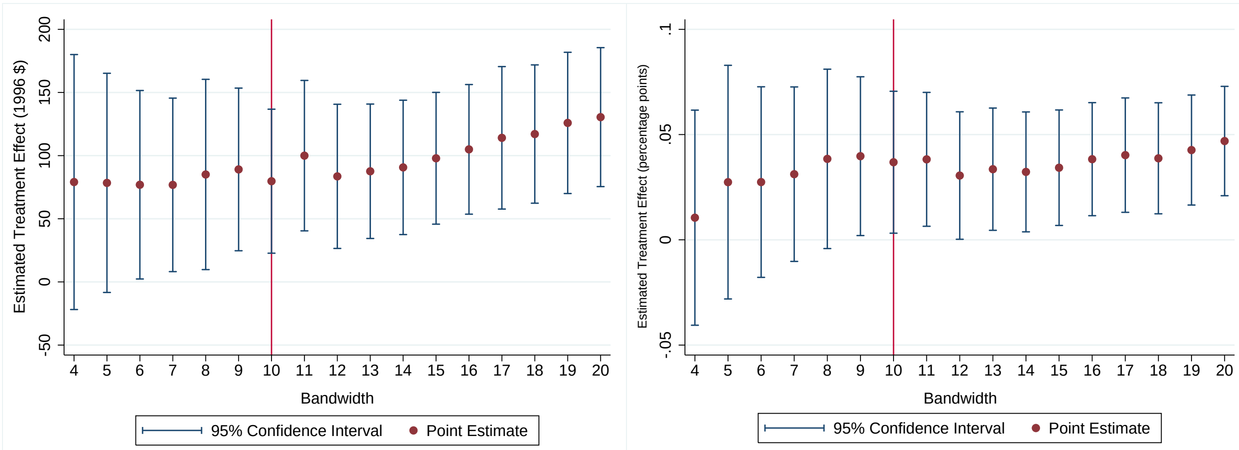

Notes: All dollar amounts are deflated to 1996 dollars using the Bureau of Labor Statistics CPI Inflation Calculator. The sample is restricted to all IRA onwers during the post reform years. Figure 5a plots estimated RDD treatment effects for the effect of eligibility for 'catch-up limits' on IRA contributions for various age window specifications. Figure 5b plots estimated RDD treatment effects for the effect of eligibility for 'catch-up limits' on the likelihood of making an IRA contribution against various age window specifications. The bandwidth of 10 used in the main results is indicated by the solid red line. The solid blue lines indicate 95% confidence intervals for the various point estimates.

the estimate of \$78 in the main results, increasing to over \$100 for the very largest bandwidth specifications. As expected, the confidence intervals narrow for wide bandwidths, although estimates are still significant at the 5 percent level for at least a bandwidth of 6. Similarly, the estimates in Figure 5b suggest that the effect of being eligible for 'catch-up limits' is not sensitive to bandwidth choice.

Table 5 examines the sensitivity of the main results to functional form choice; panels A and B of Table 5 report estimates using the full sample for IRA contributions and participation rates respectively. Panels C and D report similar estimates for the sub-sample of IRA owners. All regressions in Table 5 include year fixed effects. Each column uses a different polynomial, from a linear spline in column 1 (the baseline estimates) to a quartic spline in column 4. For brevity, the discussion below focuses on panels C and D.[45]

Coefficient estimates for β in panel C are robust to polynomials of degree two. However, precision decreases as the polynomial chosen becomes more flexible. With the cubic spline in column 3, the estimate for β falls to \$5 and is statistically indistinguishable from zero. In column 4 (which uses a quartic spline), the coefficient estimate is -\$64 and is also statistically insignificant. Inspecting the estimated level of average IRA contributions for individuals just below age 50 suggests that these higher order polynomials may not be appropriate for this application. For example, in column

[45] The results for panels A and B are similar to qualitatively to those in panels C and D.

Table 5: Sensitivity to Functional Form

Polynomial	1	2	3	4
A. Full Sample: IRA Contributions				
Control Mean	79.15	80.58	95.67	129.54
	(8.18)***	(13.87)***	(14.59)***	(14.36)***
*Over*50	22.82	22.37	4.89	-28.11
	(7.27)***	(12.81)*	(14.04)	(13.59)**
B. Full Sample: IRA Participation Rates				
Control Mean	0.040	0.041	0.052	0.071
	(0.005)***	(0.008)***	(0.007)***	(0.006)***
*Over*50	0.011	0.008	-0.005	-0.025
	(0.004)**	(0.007)	(0.007)	(0.025)***
C. IRA Owners: IRA Contributions				
Control Mean	386.97	387.74	453.02	528.07
	(33.57)***	(50.49)***	(58.63)***	(71.90)***
*Over*50	78.08	76.68	5.33	-63.93
	(28.99)**	(46.67)	(58.91)	(66.04)
D. IRA Owners: IRA Participation Rates				
Control Mean	0.193	0.198	0.243	0.295
	(0.019)***	(0.028)***	(0.027)***	(0.033)***
*Over*50	0.036	0.022	-0.029	-0.080
	(0.017)*	(0.028)	(0.028)	(0.030)**

Notes: All dollar amounts are deflated to 1996 dollars using the Bureau of Labor Statistics CPI Inflation Calculator. The dependent variable in panels A and C is a respondent's reported IRA contribution for the 2002, 2004 or 2005 calendar year. The dependent variable in panels B and D is a dummy variable equal to 1 if the respondent reported making a deductible IRA contribution for the 2002, 2004 or 2005 calendar year. *Over*50 is a dummy variable equal to 1 if the respondent is eligible for 'catch-up limits'. In all columns, the age window is 10 (respondents aged 40 to 59) and all regressions include year fixed effects. The sample in panels C and D is all IRA owners between the ages of 40 and 59. Each column corresponds to the polynomial of the assignment variable (age) in equation (2). Each observation is weighted by its inverse sampling probability in the SIPP. Across all specifications with an age window of 10, a polynomial of degree 1 minimizes the Akaike information criterion (AIC) and Bayesian information criterion (BIC). Standard errors are clustered at the age level and appear in brackets. Standard errors are clustered at the age level and appear in square brackets. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

3 the estimate for α is implies that average contributions just below age 50 are \$453. However, as shown in figure A.3b (appendix), average IRA contributions for those between the ages of 45 and 49 are never larger than \$450, and in many cases below \$400. This suggests that a cubic spline does not likely fit the data well. The same is true for the estimate of α in column 4.[46]

I repeat this analysis in panel D for IRA participation rates. Column 1 reports the specification from the main results in Table 3. The point estimate falls from 0.036 to 0.022 when a polynomial of degree 2 is chosen, though the confidence interval contain the 0.036 estimate. Inspection of figure A.4b reveals that more flexible polynomials (degree three or higher) are inappropriate for similar reasons as those mentioned in the previous paragraph.[47] Taken together, the results in

[46] Gelman and Imbens (2014) argue that higher order polynomials (cubics or higher) should not be used in RDD. They also advocate using local linear or local quadratic estimators.

[47] Across specification with an age window of 10, a polynomial of degree one (linear spline) minimizes the Akaike information criterion (AIC) and Bayesian information criterion (BIC), further that high-order polynomials likely 'over

Figure 6: Placebo Cutoffs (IRA Owners)

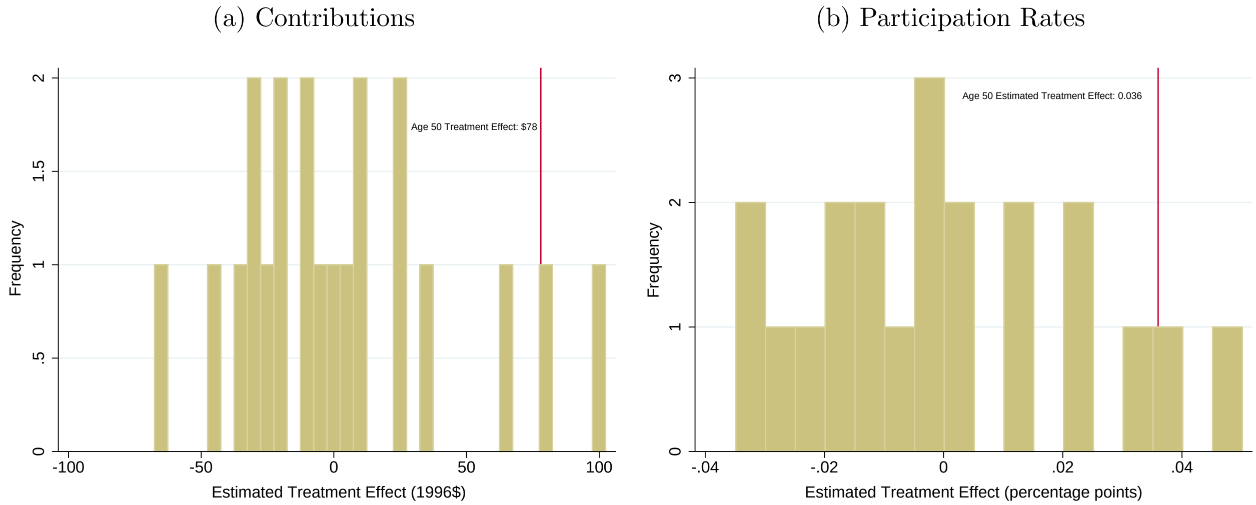

Notes: All dollar amounts are deflated to 1996 dollars using the Bureau of Labor Statistics CPI Inflation Calculator. The sample is restricted to all IRA owners during the post reform years. Figure 6a is a frequency plot of estimated RDD treatment effects for the effect of 'catch-up limits' on IRA contributions for various placebo cutoff ages. Figure 6b is a frequency plot of estimated RDD treatment effects for the effect of 'catch-up limits' on the likelihood of making an IRA contribution for various placebo age cutoffs. The solid red line indicates the estimated RDD treatment effect at the true cutoff of age 50.

Table 5 suggest that the main results are also robust to functional form choice.

It is also standard in RDD to ask whether the observed "jumps" in IRA contributions and participation rates attributed to eligibility for 'catch-up limits' are larger than or similar in magnitude to other jumps in the data. This may be especially important in this application given that savings data in survey are notoriously noisy. In Figure 6, I iteratively estimate equation (2) with year fixed effects and a bandwidth of 10 for 20 placebo age cutoffs (ages 40 to 59) and plot the frequency distribution of the resulting coefficient estimates.[48] For both IRA contributions an the likelihood of making a contribution, I find that the estimated "jump" at age 50 is the second largest of the 20 placebo RDD estimates. This is consistent with the main results reported above being statistically significant at the 5 percent level. Finally, in appendix Table A.2, I show that the main results are robust to clustering standard errors are the household level to account for possible serial correlation due to individuals who appear in the sample more than once.

5.3.2 Difference-in-Differences (DD) Estimates

The RDD estimates in Table 3 identify the effect of eligibility for catch-up limits' on IRA contributions and participation rates. However, one may also be interested in the overall evolution of

fit" the data in this application.

[48]The true cutoff of 50 is included in these 20 placebo cutoffs.

IRA contributions after the introduction of 'catch-up limits'. In this subsection, I pool the pre- and post-EGTRRA data to examine the robustness of the main results to the differences-in-differences (DD) estimator. Specifically, I estimate regressions of the form

$$R_{it} = \beta_0 + \beta_1 over50_{it} + \beta_2 post_{it} + \beta_3 over50_{it} \cdot post_{it} + \beta_4 f(age_{it}) + Z_{it}^{'}\beta_5 + e_{it} \tag{4}$$

where R_{it} is the reported IRA contribution (or participation dummy) for individual i in year t, $over50_{it}$ is a dummy variable equal to 1 if the SIPP respondent is 50 or older in year t and $post_{it}$ is a dummy variable equal to 1 in the post-EGTRRA years (2002, 2004-2005) and 0 in the pre-reform years (1996-1998, 2001). Z_{it} are baseline covariates, $f(age_{it})$ is a flexible function of age and $over50_{it} \cdot post_{it}$ is the regressor of interest.

The parameter β_2 captures the the evolution of IRA contributions after the introduction of the EGTRRA, while β_3 identifies the effect of eligibility for 'catch-up limits' (the incremental contribution limit increase for those age 50 and older). Under the standard DD assumption that the evolution of IRA contributions for those under the age of 50 is a good counterfactual for those over 50, β_3 should be identical to the RDD estimate for β. Table 6 reports the regression results for the full sample (panel A) and IRA owners (panel B). For brevity and because the unconditional (full sample) results are similar qualitatively to those for the sample of IRA owners, the discussion below focuses on panel B only.

Columns 1 and 2 of Table 6 report the DD regression results for contributions both without and with covariates while columns 3 and 4 report results from similar regressions when the dependent variable is the IRA participation dummy. Coefficient estimates for the $post_{it}$ explanatory variable are negative, suggesting that IRA contributions fell after 2002. This may reflect the volatility in equity markets in the early 2000s, making investments in financial assets less attractive for some savers. The estimate for β_3 is \$106 and is statistically significant at conventional levels. Although this estimate is approximately 30% larger than the RDD estimate in subsection 5.1., the relatively large standard errors imply that the 95% confidence interval includes the RDD estimates. The estimates for β_3 for IRA participation rates in columns 3 and 4 are similar in magnitude to the RDD estimates and are marginally insignificant at the 10 percent level. Overall, the estimates in Table 6 demonstrate that the DD estimates are similar to the main results in Table 3.

6 Crowd-Out Estimates

In this section I analyse whether increases in IRA saving lead to reductions in taxable saving. Table 7 reports the regression estimates from equation (3) for each of the taxable savings measures studied in this paper. As discussed in Section 4, I instrument IRA saving with the $over50_{it}$ dummy

Table 6: Difference-in-Differences Estimates: IRA Contributions and Participation Rates

	(1)	(2)	(3)	(4)
	Contributions		Participation Rates	
A. Full Sample = 107,596				
Post	-132.60	-332.34	-0.099	-0.201
	(202.40)	(196.87)	(0.104)	(0.102)*
*Over*50	1,001.16	1,273.86	0.157	0.285
	(833.86)	(804.58)	(0.299)	(0.296)
$Over50 * Post$	31.97	27.64	0.011	0.009
	(9.15)***	(9.05)***	(0.006)*	(0.005)
Covariates	N	Y	N	Y
B. IRA Owners = 24,488				
Post	-590.93	-704.06	-0.358	-0.362
	(743.95)	(773.05)	(0.401)	(0.409)
*Over*50	3,815.43	4,368.25	0.564	0.794
	(3,469.96)	(3,418.92)	(1.283)	(1.320)
$Over50 * Post$	106.26	102.87	0.032	0.031
	(35.61)**	(35.25)**	(0.022)	(0.022)
Covariates	N	Y	N	Y

Notes: All dollar amounts are deflated to 1996 dollars using the Bureau of Labor Statistics CPI Inflation Calculator. The dependent variable in columns 1-2 is a respondent's reported IRA contribution for the 1996-1999, 2001-2002, 2004 or 2005 calendar year. The dependent variable in columns 3-4 is a dummy variable equal to 1 if the respondent reported making a deductible IRA contribution for the 1996-1999, 2001-2002, 2004 or 2005 calendar year. *Post* is a dummy variable equal to 1 in the post EGTRRA years (2002 or later). *Over*50 is a dummy variable equal to 1 if the respondent is over the age of 50. $Over50 * Post$ is a dummy variable equal to 1 if the respondent is eligible for 'catch-up limits' (over 50 in the post- years). In all columns, the bandwidth is 10 (respondents aged 40 to 59). The sample in panel A is all respondents between the ages of 40 and 59. The sample in panel B is all IRA owners between the ages of 40 and 59. Each observation is weighted by its inverse sampling probability in the SIPP. The covariates (as defined in Table 2) are: White, Black, Hispanic, Female, Married, Personal Earned Income, Number of Kids Under the age of 18, High School, Some College and College Degree. Standard errors are clustered at the age level and appear in brackets. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

variable (equal to one if the individual is eligible for 'catch-up limits' and zero otherwise). One may be concerned that the imputed measure of IRA saving is weakly correlated with the instrument because the variation in this variable is only partially due to the higher contributions associated with 'catch-up limits'.[49] Weak instruments can lead to biased 2SLS coefficient estimates and standard errors (Murray, 2006). I test for weak instruments using the test and critical values in Stock and Yogo (2005). For the case of a single endogenous explanatory variable and a single instrument, as is the case in this application, the test statistic is the first-stage F-statistic and the critical value is 16.38. The first-stage F-statistics in Table 7 are all less than two so I cannot reject the null hypothesis that the instrument is weak. Since the 2SLS estimator is biased towards the OLS estimator when instruments are weak, I report both OLS and 2SLS estimates for equation (3) in all specifications below. The OLS and 2SLS specifications in panels A and B of Table 7 include a flexible polynomial in age and year fixed effects. This specification is similar to the specification in column 2 of Table 3.[50] Panels C and D add baseline covariates to the OLS and 2SLS specifications

[49]The remaining variation is due to accumulated interest and capital gains that are uncorrelated with eligibility for 'catch-up limits'.

[50]To be consistent with the estimates in Section 5 the regressions only include individuals within 10 years of the age 50 cutoff (40 to 59 year olds).

in panels A and B.

The 2SLS estimates for the crowd-out parameter μ_1 are generally larger in magnitude that the OLS estimates and positive in magnitude. For example, the 2SLS coefficient estimate in column 1 of panel B implies that a \$1 increase in IRA saving leads to an increase in bank saving of \$0.46, though the 95% confidence interval cannot rule out a decrease of up to \$0.28. The coefficient estimate for the change in unsecured debt is \$0.09, suggesting that a \$1 increase in IRA saving leads individuals to accumulate slightly more unsecured debt. Similar to Gelber (2011), the standard errors across all columns are large and 2SLS estimates for all measures of taxable saving are statistically indistinguishable from zero. When covariates are added, the estimates for μ_1 are similar in magnitude in panel D but remain statistically insignificant.

Interestingly, with the exception of column 5 (saving in all financial assets), the confidence intervals implied by these estimates rule out perfect (one-to-one) crowd-out. This is may not be surprising given the \$78 increase in contributions for IRA owners at age 50. If these contributions represent new saving (at least in the year the contribution is made), the implied reduction in consumption would likely be quite manageable for most individuals.[51] Table A.3 in the appendix shows the robustness of the 2SLS crowd-out estimates in Table 7 to using reported deductible IRA contributions by SIPP respondents as the endogenous variable of interest.

7 Conclusion

Using an empirical strategy that aims to address the limitations of prior research, this paper asks whether increasing IRA and 401(k) contribution limits increases personal saving. This topic is particularly relevant given the recent debates about the adequacy of retirement saving. I argue that using eligibility for 'catch-up limits,' an IRA and 401(k) contribution rule allowing individuals over the age of 50 to contribute more than those under 50, permits cleaner identification of the savings effect of these programs than previous empirical strategies. The main results indicate that eligibility for 'catch-up limits' limits leads to a large and statistically significant increase in IRA contributions. The likelihood of making a contribution also increases by 3.6 percentage points (19%) for IRA owners just over the age of 50. Furthermore, I find that average new contributions by those induced to participate are relatively small, and that eligibility for 'catch-up limits' has no significant effects on 401(k) contributions. In the standard LCM, we do not expect that a change

[51]Chetty et al. (2014) estimate relatively large crowd-out parameters using administrative data from Denmark. However, their crowd-out estimates are largest for individuals who frequently rebalance their portfolios (active savers), while crowd-out is quite small for passive individuals. The modest crowd-in (or modest crowd-out) that I find is not inconsistent with the findings in Chetty et al. (2014). If targeted advertising by banks induced those little previous active savings history to contribute (relatively small amounts) to an IRA, then much of these contributions may represent new saving.

Table 7: Crowd-Out Estimates: IRA Owners

	(1) Bank Saving	(2) Stock Saving	(3) Unsecured Debt	(4) Car Value	(5) OFA Saving
A. OLS No Covariates					
R_{it}	0.006	0.008	-0.001	0.000	0.023
	(0.005)	(0.003)**	(0.002)	(0.002)	(0.009)**
N: 6,822					
B. 2SLS No Covariates					
R_{it}	0.464	0.380	0.091	0.054	1.610
	(0.380)	(0.599)	(0.247)	(0.119)	(1.776)
F-stat (1st Stage): 1.25					
N: 6,822					
C. OLS With Covariates					
R_{it}	0.006	0.007	-0.001	0.000	0.023
	(0.005)	(0.003)***	(0.002)	(0.002)	(0.009)**
N: 6,822					
D. 2SLS With Covariates					
R_{it}	0.442	0.351	0.096	0.046	1.463
	(0.341)	(0.532)	(0.242)	(0.110)	(1.531)
F-stat (1st Stage): 1.38					
N: 6,822					

Notes: All dollar amounts are deflated to 1996 dollars using the Bureau of Labor Statistics CPI Inflation Calculator. The dependent variable is an imputed measure of taxable saving for the 2002 or 2005 calendar year (see Table 2). IRA contributions are imputed by subtracting the respondent's reported market value of their IRA account at the end of year t from the value at the end of year t - 1. In Panels B and D the level of IRA contributions (R_{it}) is instrumented with $over50_{it}$, a dummy variable equal to 1 if the respondent will be at least 50 years old by the end of the calendar year in which the interview has taken place (eligible for 'catch-up limits'). The estimation procedure is 2SLS and all regressions include year fixed effects. Each observation is weighted by its inverse sampling probability in the SIPP. The covariates are: White, Black, Hispanic, Female, Married, Personal Earned Income, Number of Kids Under the age of 18, High School, Some College and College Degree. Standard errors are clustered at the age level and appear in brackets. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

in the contribution limit will affect the choices of agents for whom the constraint is not binding. Therefore, these findings suggest that the response to eligibility for 'catch-up limits' was not limited to those constrained by the limit. They are, however, consistent with the explanation that financial institutions directed their IRA advertising towards those over the age of 50 following the introduction of 'catch-up limits', leading to higher participation rates.

Several pieces of evidence support the plausibility of this mechanism. First, descriptive evidence from financial industry publications suggest that banks and credit unions directed IRA their promotional activity towards those over the age of 50 in the years following the introduction of 'catch-up limits'. Also, since eligibility for 'catch-up limits' has insignificant effects on 401(k) contributions, it seems unlikely that actions by both employees and employer-sponsors of 401(k)s changed following the reform. Third, eligibility for 'catch-up limits' increased IRA contributions as well as the likelihood of making an contribution most for individuals with both high incomes and low levels of taxable financial assets. It is plausible that financial institutions looking to increase assets under management would target individuals in these groups. The fact that the increase in average IRA contributions is driven by jumps in participation rates at age 50 further supports this mechanism.

While informative about how expanding IRAs and 401(k)s affects contributions to these plans, the results in this paper do suffer from some limitations. First, the age-related RDD identifies the effect of eligibility for 'catch-up limits' on savings at age 50. While studying the behavior of individuals at this prime savings age is of great interest for policy makers and economists, the results in this paper cannot be interpreted as the casual effect of raising contribution limits on other age groups. Second, the findings suggest that the savings response by those eligible for 'catch-up limits' was not limited to those previously constrained by the limit. They are also consistent with an explanation that advertising by financial institutions, a potentially important determinant of retirement savings behavior, is at least partially driving the large increase in observed participation rates. However, evidence for this mechanism is suggestive and more research is needed to explore the importance of activities by third parties such as financial institutions on retirement savings behavior. Directly testing whether targeted advertising (perhaps through a field experiment) increases retirement account participation rates and contributions, as well as exploring which aspects of this advertising are most effective, represent fruitful areas for future research. Also, the results presented above do not speak to the potential dynamic responses to this reform; though recent research finds that in the domain of retirement saving, many individuals fail to act upon or understand the consequences of even contemporaneous incentives to save (Chetty et al., 2014). This suggests that any potential dynamic responses may be small. Finally, the crowd-out estimates in section 6 are imprecisely estimated, mostly because of data limitations. Drawing conclusions about the effect of raising contribution limits on total personal saving therefore requires further study.

References

Angrist, Joshua D., and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton, New Jersey: Princeton University Press.

Attanasio, Orazio P., and Thomas DeLeire. 2002. "The Effect of Individual Retirement Accounts on Household Consumption and National Saving." *The Economic Journal, 112*(481): 504-538.

Bagwell, Kyle. 2007. "The Economic Analysis of Advertising." in Mark Armstrong and Robert Porter, eds., *Handbook of Industrial Organization*, Vol. 3, Elsevier, p. 1701-1700.

Benartzi, Schlomo, and Richard Thaler. 2007. "Heuristics and Biases in Retirement Savings Behavior." *Journal of Economic Perspectives, 21*(3): 81-104

Bernheim, B. Douglas. 2002. "Taxation and Saving." in Alan J. Auerbach and Martin Feldstein, eds., *Handbook of Public Economics*, Vol. 3, Elsevier, p. 1173-1249.

Benjamin, Daniel J. 2003. "Does 401(k) eligibility increase saving? Evidence from propensity score subclassification." *Journal of Public Economics, 87*: 1259-1290.

Cameron, A. Colin, Jonah B. Gelbach and Douglas L. Miller. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *The Review of Economics and Statistics, 90*(3): 414-427.

Carroll, Gabriel D., James J. Choi, David Laibson, Brigitte C. Madrian, and Andrew Metrick. 2009. "Optimal Defaults and Active Decisions." *The Quarterly Journal of Economics, 124*(4): 1639-1674.

Chetty, Raj, John N. Friedman, Søren Leth-Petersen, Torben Heien Nielsen and Tore Olsen. 2014. "Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark." *The Quarterly Journal of Economics, 129*(3): 1141-1219.

Choi, James J., David Laibson, and Brigitte C. Madrian. 2004. "Plan Design and 401(k) Savings Outcomes." *National Tax Journal, 57*(2): 275-298.

Clark, Robert L., Jennifer A. Maki, and Melinda Sandler Morrill. 2014. "Can Simple Informational Nudges Increase Employee Participation in a 401(k) Plan?" *Southern Economic Journal, 80*(3): 677-701.

Cruz, Humberto, and Diane Lade. "Congress May Renew Efforts to Raise IRA Limits." *Sun Sentinel* March 5, 2001. Retrieved from: http://articles.sun-sentinel.com/2001-03-05/business/0103020758_1_roth-iras-retirement-savings-contribution.

Dalton, Patrick. "Tax law creates deposit potential." *ABA Bankers News* June 12, 2001. Print.

Duarte, Fabian, and Justine S. Hastings. 2012. "Fettered Consumers and Sophisticated Firms: Evidence from Mexico's Privatized Social Security Market." NBER Working Paper 18582, National Bureau of Economic Research.

Duflo, Esther, William Gale, Jeffrey Liebman, Peter Orszag, and Emmanuel Saez. 2006. "Saving Incentives for Low- and Middle-Income Families: Evidence from a Field Experiment with H&R Block." *The Quarterly Journal of Economics, 121*(4): 1311-1346.

Duflo, Esther, William Gale, Jeffrey Liebman, Peter Orszag, and Emmanuel Saez. 2007. "Savings Incentives for Low- and Middle-Income Families in the United States: Why is the Saver's Credit not More Effective?" *Journal of the European Economic Association, 5*(2-3): 647-661.

Ellens, Jennifer. "Retirement Catch-up Contributions." *ABA Trust & Investments* January 2004. Print.

Engelhardt, Gary V., and Anil Kumar. 2007. "Employer matching and 401(k) saving: Evidence from the health and retirement study." *Journal of Public Economics, 91*(10): 1920-1943.

Feenberg, Daniel, and Jonathan Skinner. 1989. "Sources of IRA savings." in L.H. Summers, ed., *Tax Policy and the Economy*, Vol. 3 of NBER Chapters, Cambridge, Mass.: National Bureau of Economic Research, p. 25-46.

Gale, William G., and John Karl Scholz. 1994. "IRAs and Household Saving." *The American Economic Review, 84*(5): 1233-1260.

Gelber, Alexander M. 2011. "How do 401(k)s Affect Saving? Evidence from Changes in 401(k) Eligibility." *American Economic Journal: Economic Policy, 3*(1): 103-122.

Gelman, Andrew and Guido Imbens. 2014. "Why Higher-Order Polynomials Should Not Be Used in Regression Discontinuity Designs." NBER Working Paper 20405, National Bureau of Economic Research.

Gurun, Umit G., Gregot Matvos, and Amit Seru. 2013. "Advertising Expensive Mortgages." NBER Working Paper 18910, National Bureau of Economic Research.

Hastings, Justine S., Ali Hortaçsu, and Chad Syverson. "Advertising and Competition in Privatized Social Security: The Case of Mexico." NBER Working Paper 18881, National Bureau of Economic Research.

"Help with Retirement Planning." *Teller Sense* January 15, 2003. Print.

Holden, Sarah, and Steven Bass. 2014. "The IRA Investor Profile: Traditional IRA Investors' Activity, 2007-2012." Investment Company Institute (ICI) Research Report.

Holden, Sarah, and Daniel Schrass. 2014. 'Defined Contribution Plan Participants' Activities, 2013." Investment Company Institute (ICI) Research Report.

Honka, Elisabeth, Ali Hortaçsu, and Maria Ana Vitorino. 2014. "Advertising, Consumer Awareness and Choice: Evidence from the U.S. Banking Industry." Unpublished manuscript.

Imbens, Guido W., and Karthik Kalyanaraman. 2012. "Optimal Bandwidth Choice for the Regression Discontinuity Estimator." *Review of Economic Studies, 79*(3): 933-959.

Investment Company Institute. 2012. "IRAs in Americans' Retirement Preparedness." Investment Company Institute (ICI) Research Report.

Karlan, Dean, Margaret McConnell, Sendhil Mullainathan, and Jonathan Zinman. 2010. "Getting to the top of Mind: How Reminders Increase Saving." NBER Working Paper 16205, National Bureau of Economic Research.

Lee, David S., and David Card. 2008. "Regression Discontinuity Inference with Specification Error." *Journal of Econometrics, 142*(2): 655-674.

Lee, David S., and Thomas Lemieux. 2010. "Regression Discontinuity Designs in Economics." *Journal of Economic Literature, 48*(2): 281-355.

Madrian, Brigitte C., and Dennis F. Shea. 2001. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *The Quarterly Journal of Economics, 116*(4): 1149-1187.

Milligan, Kevin. 2003. "How do contribution limits affect contributions to tax-preferred savings accounts?" *Journal of Public Economics, 87*(2): 253-281.

Murray, Michael P. 2006. "Avoiding Invalid Instruments and Coping with Weak Instruments." *Journal of Economic Perspectives, 20*(4): 111-132.

Orszag, Peter, and Orszag, Jonathan. 2001. "Would Raising IRA Contribution Limits Bolster Retirement Security for Lower- and Middle-Income Families?" Center on Budget and Policy Priorities, Washington D.C.

Poterba, James M. 2014. "Retirement Security in an Aging Society." NBER Working Paper 19930, National Bureau of Economic Research.

Poterba, James M., Steven F. Venti, and David A. Wise. 1995. "Do 401(k) contributions crowd-out other personal saving?" *Journal of Public Economics, 58*(1): 1-32.

Ramnath, Shanthi. 2013. "Taxpayers' responses to tax-based incentives for retirement savings: Evidence from the Saver's Credit notch." *Journal of Public Economics, 101*: 77-93.

Saez, Emmanuel. 2009. "Details Matter: The Impact of Presentation and Information on the Take-up of Financial Incentives for Retirement Saving." *American Economic Journal: Economic Policy, 1*(1): 204-228.

Stock, James H., and Motohiro Yogo. 2005. "Testing for Weak Instruments in Linear IV Regression." in D. Andrews, ed., *Identification and Inference for Econometric Models*. New York: Cambridge University Press.

Venti, Steven F., and David A. Wise. 1990a. "Have IRAs Increased US Saving? Evidence from Consumer Expenditure Surveys." *The Quarterly Journal of Economics, 105*(3): 661-698.

Venti, Steven F., and David A. Wise. 1990b. "The Determinants of IRA Contributions and the Effect of Limit Changes." NBER Working Paper 1731, National Bureau of Economic Research.

Zuehlke, Dennis. 2001. "Tax-cut bill creates new IRA opportunities." *Credit Union Executive Journal, 41*: 12-16.

Appendix

A.1

The definitions of the imputed taxable savings measures used in this paper are listed below:

- **Bank Savings:** changes in the interest earning assets held in banking institutions (such as interest earning savings accounts and certificates of deposit
- **Stock Savings:** changes in equity held in stocks and mutual fund shares
- **Unsecured Debt:** change in total unsecured debt (such as credit card and unsecured lines of credit)
- **Car Value:** changes in the net equity of vehicles
- **OFA Savings:** changes in the sum of total financial assets (bank saving, stock saving, interest earning assets not in financial institutions and equity in other assets)

A.2

In this appendix section, I first describe the Kuhn-Tucker cases that outline the 5 possible optimal choices for R_{1i} and N_{1i}. The proof for Proposition 1 follows.

Optimization Problem

$$\max_{R_{1i},N_{1i}} \; u(c_{1i}, c_{2i}) - v(R_{1i})$$

subject to:

$$c_{1i} = \begin{cases} (y - R_{1i})(1-\tau_1) - N_{1i} - k_i & \text{if } R_{1i} > 0 \\ y(1-\tau_1) - N_{1i} & \text{if } R_{1i} = 0 \end{cases}$$

$$c_{2i} = (1+r)(1-\tau_2)R_{1i} + (1 + r(1-\tau_2))N_{1i}$$

$$R_{1i} \geq 0; \quad R_{1i} \leq L; \quad N_{1i} \geq 0$$

Let ϕ and μ be the Lagrange multipliers for the non-negativity constraints for the retirement account and taxable account respectively. Also, let λ be the Lagrange multiplier for the contribution limit constraint.

First Order Conditions:

$$R_{1i}: \quad -u_1(1-\tau_1) + u_2(1+r)(1-\tau_2) - v'(R_{1i}) + \phi - \lambda = 0$$

$$N_{1i}: \quad -u_1 + u_2(1 + r(1-\tau_2)) + \mu = 0$$

Kuhn-Tucker Conditions:

$$\phi \geq 0; \quad R_{1i} \geq 0; \quad \phi R_{1i} = 0$$

$$\lambda \geq 0; \quad R_{1i} \leq L; \quad \lambda(R_{1i} - L) = 0$$

$$\mu \geq 0; \quad N_{1i} \geq 0; \quad \mu N_{1i} = 0$$

Let (R^*_{1i}, N^*_{1i}) be the solution to this optimization problem for individual i. Since $L > 0$, $R^*_{1i} = L$

and $R_{1i} = 0$ cannot both be observed simultaneously. This reduces the number of possible Kuhn-Tucker cases from 8 to 6.

Case A: $R^*_{1i} = 0$ **AND** $N^*_{1i} = 0$
This case is impossible because it implies that $S^*_i = R^*_{1i} + N^*_{1i} = 0$ so there is no period 2 consumption. Assuming $u_2 > 0$ and diminishing marginal utility, this case is not a solution candidate.

Case B: $R^*_{1i} = 0$ **AND** $N^*_{1i} > 0$
This implies that $\lambda^* = 0$, $\mu^* = 0$ and $\phi^* \geq 0$. From the first order conditions we have:

$$u_1 = u_2(1 + r(1 - \tau_2))$$

$$\phi^* = u_1(1 - \tau_1) - u_2(1 + r)(1 - \tau_2) + v'(0)$$

For $\phi^* \geq 0$ we require that $v'(0) \geq u_2(\tau_1 - \tau_2 + r\tau_1(1 - \tau_2))$.

Case C: $0 < R^*_{1i} < L$ **AND** $N^*_{1i} = 0$
This implies that $\lambda^* = 0$, $\phi^* = 0$ and $\mu^* \geq 0$. From the first order conditions we have:

$$v'(R^*_{1i}) = -u_1(1 - \tau_1) + u_2(1 + r)(1 - \tau_2)$$

$$\mu^* = \frac{u_2(\tau_1 - \tau_2 + r\tau_1(1 - \tau_2)) - v'(R^*_{1i})}{1 - \tau_1}$$

The first equation pints down R^*_{1i}. For $\mu^* \geq 0$ we require that $v'(R^*_{1i}) \leq u_2(\tau_1 - \tau_2 + r\tau_1(1 - \tau_2))$.

Case D: $0 < R^*_{1i} < L$ **AND** $N^*_{1i} > 0$
This implies that $\lambda^* = 0$, $\phi^* = 0$ and $\mu^* = 0$. From the first order conditions we have:

$$v'(R^*_{1i}) = -u_1(1 - \tau_1) + u_2(1 + r)(1 - \tau_2)$$

$$u_1 = u_2(1 + r(1 - \tau_2))$$

These equations pin down R^*_{1i} and N^*_{1i} subject to $0 \leq R^*_{1i} \leq L$ AND $N^*_{1i} \geq 0$. These equations can also be rearranged to show that the individual determines their optimal retirement account and taxable account contributions by comparing the relative disutility associated with incremental contributions to the retirement account (due to the liquidity penalty) and the period 2 consumption benefit associated with higher contributions (i.e. period 2 marginal utility plus the net price incentive from contributing). Formally, we have:

$$v'(R^*_{1i}) = u_2(\tau_1 - \tau_2 + r\tau_1(1 - \tau_2))$$

Case E: $R_{1i}^* = L$ **AND** $N_{1i}^* = 0$
This implies that $\phi^* = 0$, $\lambda^* \geq 0$ and $\mu^* \geq 0$. From the first order conditions we have:

$$\lambda^* = u_2(1+r)(1-\tau_2) - u_1(1-\tau_1) - v'(L)$$

$$\mu^* = u_1 - u_2(1 + r(1-\tau_2))$$

For $\lambda^* \geq 0$ we require that $v'(L) \leq -u_1(1-\tau_1) + u_2(1+r)(1-\tau_2)$. For $\mu^* \geq 0$ we require that $u_1 \geq u_2(1+r(1-\tau_2))$. This last inequality says that for $\mu^* \geq 0$ we require that the MRS between period 1 and period 2 consumption to be at least as large as the net price benefit of contributing to the taxable account.

Case F: $R_{1i}^* = L$ **AND** $N_{1i}^* > 0$
This implies that $\phi^* = 0$, $\lambda^* \geq 0$ and $\mu^* = 0$. From the first order conditions we have:

$$\lambda^* = u_2(1+r)(1-\tau_2) - u_1(1-\tau_1) - v'(L)$$

$$u_1 = u_2(1 + r(1-\tau_2))$$

For $\lambda^* \geq 0$ we require that $v'(L) \leq u_2(\tau_1 - \tau_2 + r\tau_1(1-\tau_2))$.

Impact of Fixed Costs
Due to the fixed cost for retirement account contributors, positive contributions must also provide a higher level of utility than not contributing. Specifically, in addition to satisfying the relevant Kuhn-Tucker conditions, any solution for which positive retirement account contributions is optimal (Cases C, D, E and F) with solution pair ($R_{1i}^* \in (0, L]$, N_{1i}^*) must satisfy for k_i and any $\tilde{N}_{1i}$:

$$u(R_{1i}^*, N_{1i}^*; k_i) - v(R_{1i}^*) \geq u(0, \tilde{N}_{1i}; 0)$$

Proposition 1: *In the standard LCM, an increase in the contribution limit L: (a) increases retirement account contributions only for type l savers that are limit contributors, (n) decreases taxable account contributions only for type l savers that are limit contributors and (c) increases total personal saving for type l savers that are limit contributors.*

Proof: Consider each part of the proposition in turn:

(a) Increases retirement account contributions only for type l savers that are limit contributors: Recall that type m and type h savers do not contribute to the retirement account because of high fixed costs. For some type l savers in Cases B, C and D it is optimal to contribute less than the limit to the retirement account. Irrespective of fixed costs, there is no effect of an increase in L for these individuals. For type l savers for whom the initial limit is binding ($R^*_{1i} = L$), an increase in the contribution limit leads to higher retirement account contributions:

$$\frac{\partial R^*_{1i}}{\partial L} = 1 > 0$$

(b) Decreases taxable account contributions only for type l savers that are limit contributors: Since non-limit contributors are not impacted by an increase in L, taxable account contributions are unaffected for these individuals. For type l individuals in Case E, $N^*_{1i} = 0$ so that $\frac{\partial N^*_{1i}}{\partial L} = 0$.

For type l individuals in Case F, N^*_{1i} is determined by $u_1 = u_2(1 + r(1 - \tau_2))$. Differentiating this equation with respect to L, substituting $\frac{\partial R^*_{1i}}{\partial L} = 1$ and rearranging yields:

$$\frac{\partial N^*_{1i}}{\partial L} = \frac{u_{11}(1-\tau_1) - u_{12}((1+r)(1-\tau_2) + (1+r(1-\tau_2))(1-\tau_1)) + u_{22}(1+r(1-\tau_2))(1-\tau_2)}{-u_{11} + 2u_{12}(1+r(1-\tau_2)) - u_{22}(1+r(1-\tau_2))^2} < 0$$

Therefore, an increase in L decreases taxable account contributions for type l limit contributors.

(c) Increases the total personal saving for type l savers that are limit contributors: Recall that total personal saving is defined as $S^*_i = R^*_{1i} + N^*_{1i}$. Differentiating this expression with respect to L, given parts (a) and (b) yields:

$$\frac{\partial S^*_i}{\partial L} = \frac{\partial R^*_{1i}}{\partial L} + \frac{\partial N^*_{1i}}{\partial L} = \frac{-u_{11}\tau_1 + u_{12}(\tau_1 - \tau_2 + r\tau_1(1-\tau_2)) - u_{22}(1+r(1-\tau_2))\tau_2}{-u_{11} + 2u_{12}(1+r(1-\tau_2)) - u_{22}(1+r(1-\tau_2))^2} > 0$$

Since type l non-limit contributors are not affected by an increase in L and limit contributors in Case E do not experience a change in taxable account contributions, the expression above confirms that total personal saving increases. $\square$

A.3

In this appendix, I derive the formula on page 26. Recall that the increase in (unconditional) average IRA contributions due to the 'catch-up limit' can be expressed as

$$\beta^C = \underbrace{\big(P[R_{it} > 0|D_{it} = 1] - P[R_{it} > 0|D_{it} = 0]\big)}_{\text{Participation Effect}} E[R_{it}|R_{it} > 0, D_{it} = 1]$$
$$+ \underbrace{\big(E[R_{it}|R_{it} > 0, D_{it} = 1] - E[R_{it}|R_{it} > 0, D_{it} = 0]\big)}_{\textit{Conditional-on-positive (COP) Effect}} P[R_{it} > 0|D_{it} = 0]$$

Let $\beta^P = P[R_{it} > 0|D_{it} = 1] - P[R_{it} > 0|D_{it} = 0]$ denote the participation effect and let the the expectation $\alpha = P[R_{it} > 0|D_{it} = 0]$ denote the probability of making an IRA contribution for those just below age 50. Also, let $s = M$ denote those (marginal savers) induced by 'catch-up limits' to participate in IRAs. Let $s = I$ denote the inframarginal savers who would have contributed to an IRA in the absence of the reform. The fraction of IRA contributors that are "new contributors" is $\psi = \frac{\beta^P}{\alpha+\beta^P}$. Therefore, the expression above can be rewritten as

$$\beta^C = \beta^P\Big(\psi E[R_{it}|R_{it} > 0, D_{it} = 1, s = M] + (1-\psi)E[R_{it}|R_{it} > 0, D_{it} = 1, s = I]\Big)$$
$$+ \Big(\psi E[R_{it}|R_{it} > 0, D_{it} = 1, s = M] + (1-\psi)E[R_{it}|R_{it} > 0, D_{it} = 1, s = I] - E[R_{it}|R_{it} > 0, D_{it} = 0]\Big)\alpha$$

Collecting terms and rearranging yields

$$\beta^C = \psi E[R_{it}|R_{it} > 0, D_{it} = 1, s = M]\Big(\alpha + \beta^P\Big) + (1-\psi)E[R_{it}|R_{it} > 0, D_{it} = 1, s = I]\Big(\alpha + \beta^P\Big)$$
$$- \alpha E[R_{it}|R_{it} > 0, D_{it} = 0]$$

Since $\psi = \frac{\beta^P}{\alpha+\beta^P}$, $\psi(\alpha + \beta^P) = \beta^P$ and $(1-\psi)(\alpha + \beta^P) = \alpha$. Therefore, the expression above can be rewritten as

$$\beta^C = \beta^P E[R_{it}|R_{it} > 0, D_{it} = 1, s = M] + \alpha\Big(E[R_{it}|R_{it} > 0, D_{it} = 1, s = I] - E[R_{it}|R_{it} > 0, D_{it} = 0]\Big)$$

Figure A.1: Distribution of the Assignment Variable

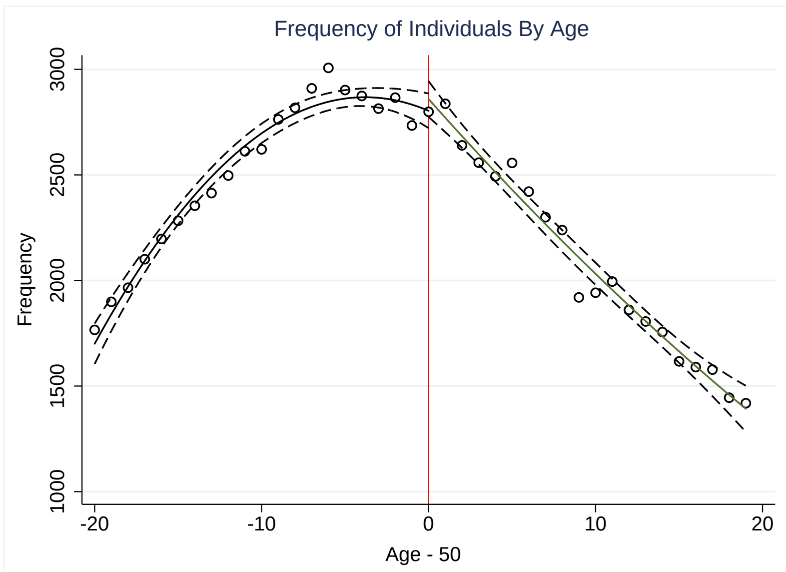

Notes: This figure plots the number of observations in each age bin the post-reform (EGTRRA 2001) SIPP data for individuals between the age of 40 and 59. Observations are from wave 7 of the 2001 SIPP panel (2002 calendar year) and waves 4 and 7 of the 2004 SIPP panel (2004 and 2005 calendar years). The solid red line refers to the normalized age 50 cutoff above which individuals are eligible for 'catch-up limits'. The solid black lines represent the fitted values from quadratic splines (where the dependent variable is the number of observations in each age bin) and standard errors are clustered at the age level. The regression weights each age bin frequency by it's relative weight for the total population. The corresponding 95 percent confidence intervals are represented by the dashed lines. The estimate for the discontinuity at the age 50 cutoff is 35.02 (s.e. = 73.09).

Figure A.2: Covariates Balance Test

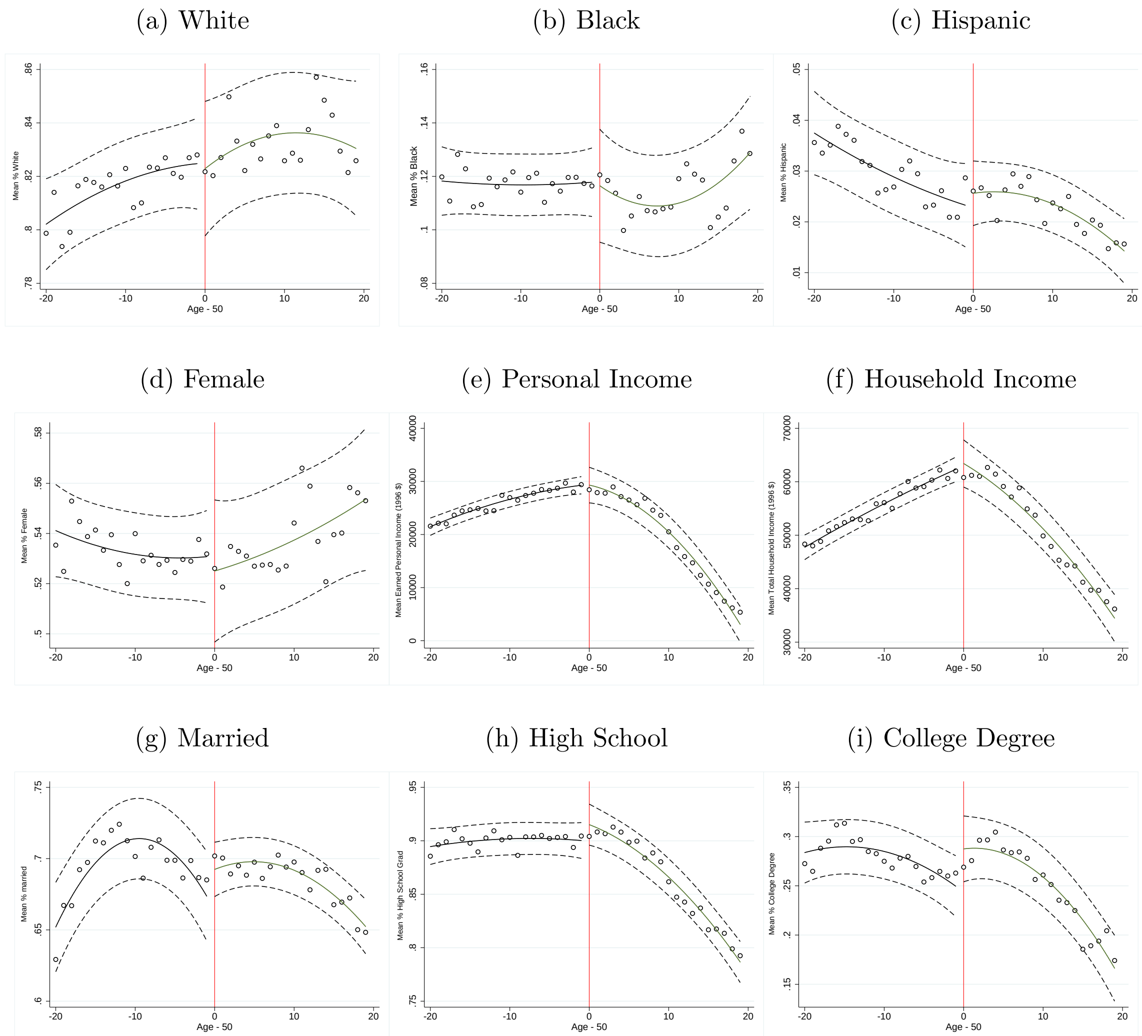

Notes: This figure plots covariates against age for the same sample as in Figure A.1. The solid red line refers to the normalized age 50 cutoff above which individuals are eligible for 'catch-up limits'. The solid black lines represent the fitted values from quadratic splines (where the dependent variable is the relevant covariate). The corresponding 95 percent confidence intervals are represented by the dashed lines.

Figure A.3: Mean IRA Contributions by Age: IRA Owners

(a) Pre-EGTRRA

(b) Post-EGTRRA

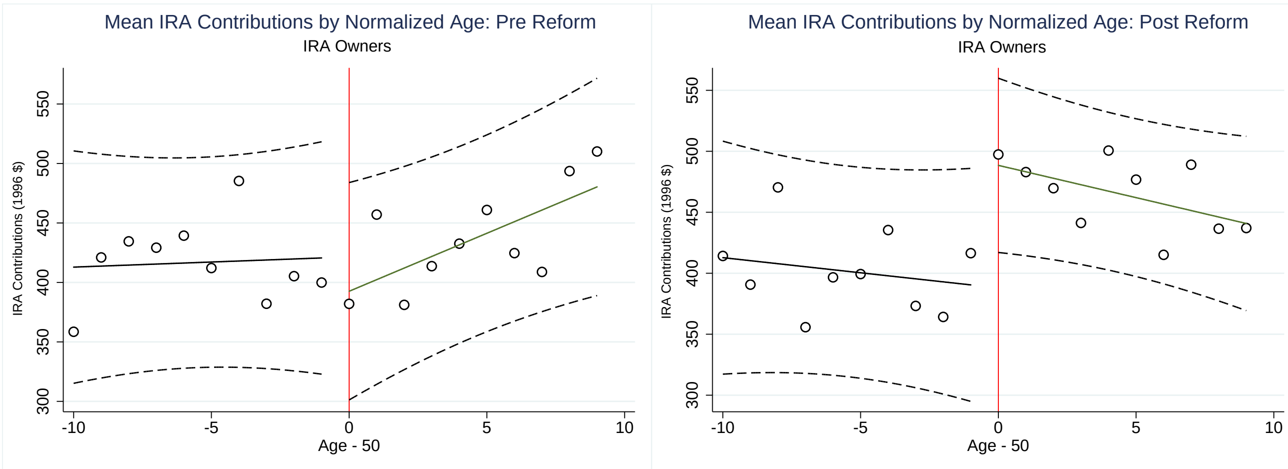

Notes: All dollar amounts are deflated to 1996 dollars using the Bureau of Labor Statistics CPI Inflation Calculator. The sample is restricted to all individuals between the ages of 40 and 59 during the pre and post-reform years. Figure A.3a plots average IRA contributions against the normalized assignment variable (age) for the 1996-1998 and 2001 pre-reform years for IRA owners. Figure A.3b plots average IRA contributions against the normalized assignment variable (age) for the 2002 and 2004-2005 post-reform years for IRA owners. Linear spline regressions are estimated on each side of the cutoff and the corresponding 95 percent confidence intervals are represented by the dashed lines.

Figure A.4: Mean IRA Participation Rates by Age: IRA Owners

(a) Pre-EGTRRA

(b) Post-EGTRRA

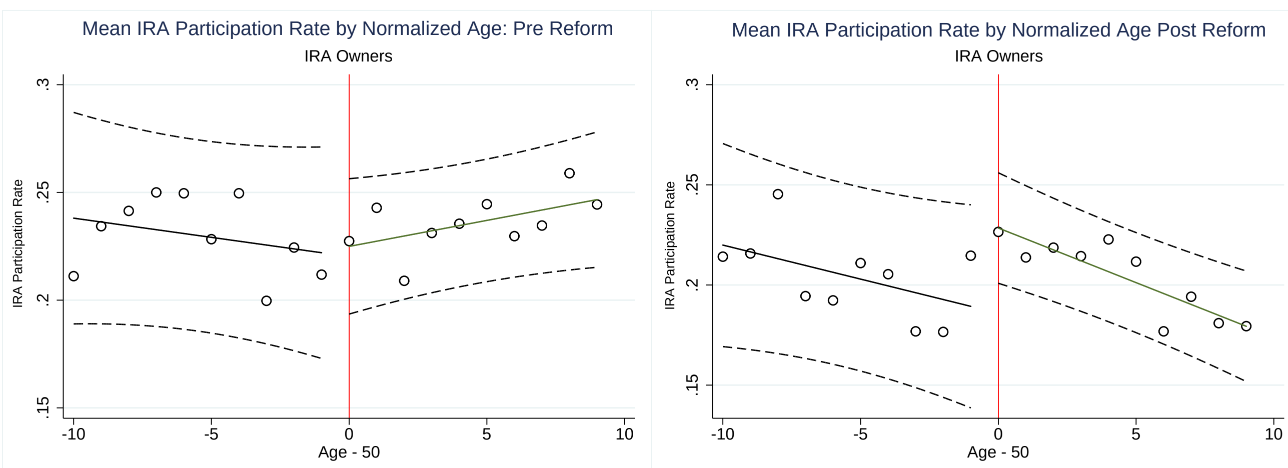

Notes: The sample is restricted to all individuals between the ages of 40 and 59 during the pre and post-reform years. Figure A.4a plots IRA participation rates against the normalized assignment variable (age) for the 1996-1998 and 2001 pre-reform years for IRA owners. Figure A.4b plots IRA participation rates against the normalized assignment variable (age) for the 2002 and 2004-2005 post-reform years for IRA owners. Linear spline regressions are estimated on each side of the cutoff and the corresponding 95 percent confidence intervals are represented by the dashed lines.

Figure A.5: Mean 401(k) Contributions by Age

(a) Pre-EGTRRA (b) Post-EGTRRA

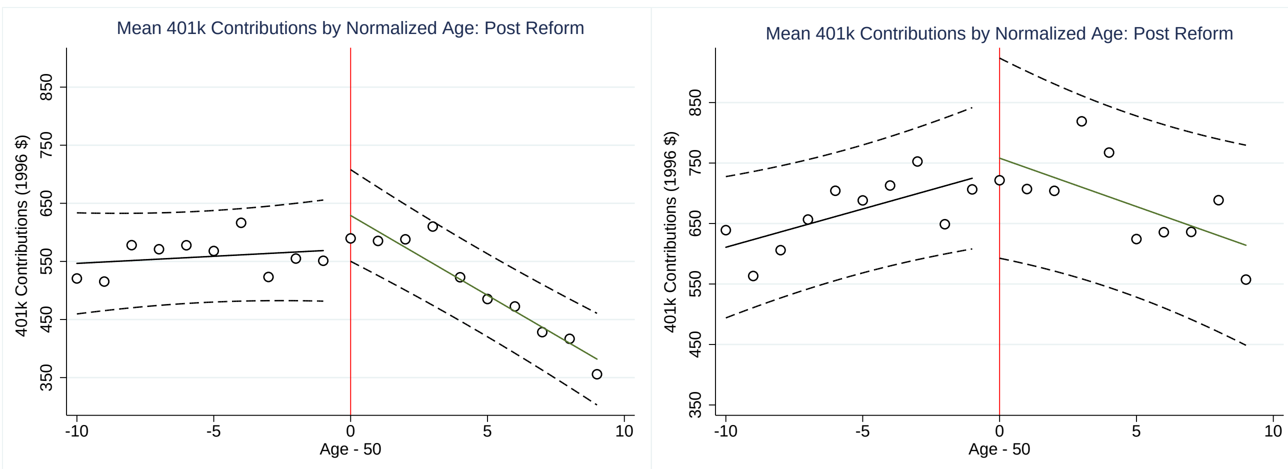

Notes: All dollar amounts are deflated to 1996 dollars using the Bureau of Labor Statistics CPI Inflation Calculator. The sample is restricted to all individuals between the ages of 40 and 59 during the pre and post-reform years. Figure A.5a plots (unconditional mean) 401(k) contributions against the normalized assignment variable (age) for the 1996-1998 and 2001 pre-reform years. Figure A.5b plots (unconditional mean) 401(k) contributions against the normalized assignment variable (age) for the 2002 and 2004-2005 post-reform years. Linear spline regressions are estimated on each side of the cutoff and the corresponding 95 percent confidence intervals are represented by the dashed lines.

Table A.1: Covariate Balance Test

	Full Sample (N=54,771)		IRA Owners(N=12,261)	
Covariate	Control Mean	$Over50$	Control Mean	$Over50$
White	0.844	-0.001	0.937	-0.007
	(0.004)***	(0.005)	(0.008)***	(0.008)
Black	0.112	-0.002	0.032	-0.005
	(0.003)***	(0.003)	(0.006)***	(0.006)
Hispanic	0.008	0.003	-0.002	0.005
	(0.003)***	(0.003)	(0.004)	(0.005)
Female	0.521	-0.003	0.539	-0.020
	(0.003)***	(0.004)	(0.013)***	(0.016)
Married	0.698	0.006	0.738	-0.007
	(0.007)***	(0.007)	(0.013)***	(0.017)
Personal Earned Income	29,695	-467	43,144	1
	(685)***	(668)	(1,460)***	(1,086)
Household Income	62,348	-187	83,410	1,580
	(1,032)***	(1,119)	(1,808)***	(1,720)
High School	0.294	-0.008	0.190	-0.002
	(0.014)***	(0.009)	(0.015)***	(0.015)
Some College	0.320	-0.011	0.281	-0.006
	(0.007)***	(0.006)*	(0.013)***	(0.017)
College Degree	0.264	0.030	0.505	0.014
	(0.005)***	(0.009)***	(0.015)***	(0.021)
Veteran Status	0.093	0.004	0.063	0.015
	(0.005)***	(0.006)	(0.011)***	(0.014)
Prob > χ^2	0.4010		0.6456	

Notes: The dependent variables are: White, Black, Hispanic, Female, Married, Personal Earned Income, Household Income, High School (highest degree), Some College (highest degree), College Degree (highest degree) and Veteran Status. $Over50$ is a dummy variable equal to 1 if the respondent will be at least 50 years old by the end of the calendar year in which the interview has taken place (eligible for 'catch-up limits'). Linear splines (with an age window of 10) are estimated and all regressions have year fixed effects. Each observation is weighted by its inverse sampling probability in the SIPP. The bottom row reports p-values from a Wald Test of joint-significance of the coefficients on the $Over50$ dummies, except College Degree. When College Degree is included in the Seemingly Unrelated Regression, the p-values are 0.01 for the Full Sample and 0.550 for IRA owners. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A.2 Robustness: Effect of ‘Catch-up Limits’ on IRA Contributions and Participation Rates

	(1) Baseline	(2) Cluster HH se	(3) Robust se	(4) Bootstrap se
A. Full Sample: IRA Contributions (N= 51,725)				
Control Mean	79.15	79.15	79.15	72.48
	(8.18)***	(9.44)***	(8.31)***	(5.54)***
$Over50$	22.82	22.82	22.82	26.75
	(7.27)***	(11.11)**	(10.23)**	(5.24)***
B. Full Sample: IRA Participation (N = 51,725)				
Control Mean	0.040	0.040	0.040	0.037
	(0.005)***	(0.004)***	(0.003)***	(0.004)***
$Over50$	0.011	0.011	0.011	0.011
	(0.004)**	(0.005)**	(0.004)**	(0.003)***
C. IRA Owners: Contributions (N = 12,261)				
Control Mean	386.97	386.97	386.97	368.11
	(33.57)***	(40.28)***	(35.87)***	(26.69)***
$Over50$	78.08	78.08	78.08	101.14
	(28.99)**	(43.53)*	(40.30)*	(25.18)***
D. IRA Owners: Participation (N = 12,261)				
Control Mean	0.193	0.193	0.193	0.185
	(0.019)***	(0.016)***	(0.014)***	(0.016)***
$Over50$	0.036	0.036	0.036	0.042
	(0.017)*	(0.017)**	(0.016)**	(0.014)***

Notes: All dollar amounts are deflated to 1996 dollars using the Bureau of Labor Statistics CPI Inflation Calculator. Standard errors for point estimates appear in brackets. In Panels A and C the dependent variable is a respondent's reported IRA contribution for the 2002, 2004 or 2005 calendar year. In Panels B and D the dependent variable is a dummy variable equal to 1 if the respondent made an IRA contribution in the given year. $Over50$ is a dummy variable equal to 1 if the respondent will be at least 50 years old by the end of the calendar year in which the interview has taken place (eligible for ‘catch-up limits’). Column 1 reports the baseline estimates from Table 3. In column 2, standard errors are clustered at the household level. In column 3, heteroskedasticity robust standard errors are reported. In column 4, bootstrap standard errors are clustered at the birth-cohort (age) level with 100 replications. Except for the estimates in column 4, each observation is weighted by its inverse sampling probability in the SIPP. Linear splines are estimated on each side of the cutoff and all regressions include year fixed effects. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A.3 Alternate Crowd-Out Estimates: IRA Owners

	(1) Bank Saving	(2) Stock Saving	(3) Unsecured Debt	(4) Car Value	(5) OFA Saving
A. OLS No Covariates					
R_{it}	0.125	0.195	-0.077	0.031	0.743
	(0.241)	(0.127)	(0.074)	(0.106)	(0.548)
N: 6,822					
B. 2SLS No Covariates					
R_{it}	14.35	11.35	2.89	1.71	50.54
	(14.75)	(17.09)	(7.05)	(3.54)	(45.69)
F-stat (1st Stage): 3.60					
N: 6,822					
C. OLS With Covariates					
R_{it}	0.137	0.204	-0.083	0.027	0.775
	(0.245)	(0.125)	(0.073)	(0.105)	(0.546)
N: 6,822					
D. 2SLS With Covariates					
R_{it}	17.27	13.36	3.78	1.87	58.42
	(18.55)	(20.81)	(8.74)	(4.09)	(56.80)
F-stat (1st Stage): 2.82					
N: 6,822					

Notes: All dollar amounts are deflated to 1996 dollars using the Bureau of Labor Statistics CPI Inflation Calculator. The dependent variable is an imputed measure of taxable saving for the 2002 or 2005 calendar year (see Table 2). In panels B and D the level of reported deductible IRA contributions (R_{it}) is instrumented with $over50_{it}$, a dummy variable equal to 1 if the respondent will be at least 50 years old by the end of the calendar year in which the interview has taken place (eligible for 'catch-up limits'). The estimation procedure is 2SLS and all regressions include year fixed effects. Each observation is weighted by its inverse sampling probability in the SIPP. The covariates are: White, Black, Hispanic, Female, Married, Personal Earned Income, Number of Kids Under the age of 18, High School, Some College and College Degree. Standard errors are clustered at the age level and appear in brackets. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A.4 Effect of ‘Catch-up Limits’: OLS (unweighted) Estimates

	(1)	(2)	(3)	(4)	(5)	(6)
		Contributions			Participation Rates	
A. Full Sample = 51,725						
Control Mean	90.47	72.48	-72.84	0.044	0.037	-0.023
	(4.78)***	(4.69)***	(4.96)***	(0.003)***	(0.004)***	(0.006)***
$Over50$	26.62	26.75	23.35	0.011	0.011	0.010
	(4.78)***	(4.69)***	(4.95)***	(0.003)***	(0.003)***	(0.003)***
Year FE	N	Y	Y	N	Y	Y
Covariates	N	N	Y	N	Y	Y
B. IRA Owners = 12,261						
Control Mean	387.82	368.11	184.77	0.186	0.185	0.203
	(18.36)***	(25.71)***	(78.49)***	(0.013)***	(0.016)***	(0.036)***
$Over50$	100.54	101.14	93.71	0.042	0.042	0.040
	(20.20)***	(20.48)***	(19.59)***	(0.014)***	(0.014)***	(0.013)***
Year FE	N	Y	Y	N	Y	Y
Covariates	N	N	Y	N	N	Y
C. 401(k) Full Sample: N=29,365						
Control Mean	1,151.98	749.02	-654.26	0.327	0.233	-0.049
	(39.97)***	(47.01)***	(70.17)***	(0.008)***	(0.010)***	(0.014)***
$Over50$	34.63	36.33	-9.21	-0.009	-0.008	-0.016
	(52.30)	(50.60)	(41.78)	(0.011)	(0.011)	(0.009)
Year FE	N	Y	Y	N	Y	Y
Covariates	N	N	Y	N	N	Y

Notes: All dollar amounts are deflated to 1996 dollars using the Bureau of Labor Statistics CPI Inflation Calculator. The dependent variable in columns 1-3 is a respondent’s reported IRA or 401(k) contribution for the 2002, 2004 or 2005 calendar year. The dependent variable in columns 4-6 is a dummy variable equal to 1 if the respondent reported making a deductible IRA or 401(k) contribution for the 2002, 2004 or 2005 calendar year. $Over50$ is a dummy variable equal to 1 if the respondent is eligible for ‘catch-up limits’. In all columns, the age window 10 on either side of the cutoff (respondents aged 40 to 59) and regression results correspond to linear splines around the age 50 cutoff. The sample in panel A is all IRA respondents between the ages of 40 and 59. The sample in panel B is all IRA owners between the ages of 40 and 59. The sample in panel C is all 401(k) respondents between the ages of 40 and 59. The covariates (as defined in Table 2) are: White, Black, Hispanic, Female, Married, Personal Earned Income, Number of Kids Under the age of 18, High School, Some College and College Degree. Standard errors are clustered at the age level and appear in brackets. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.