

Characteristics and Motives of Roth Converters: Evidence from the 2010 Repeal of Income Limits

Lucas Goodman* and Heidi R. Schramm*

May 2018

Abstract

We study Roth conversions using administrative tax data on the universe of individuals who made conversions from 2003 through 2015. Specifically, we document time trends in conversions, investigate how the 2010 removal of Roth conversion income limits affected conversions in that year, show the extent to which removal of these limits allowed for “backdoor Roth” contributions thereafter, and explore potential motives for Roth conversions. We find aggregate conversions increased by over 700% in 2010 compared to the prior year and that an increasing share of converters are backdoor converters: while approximately 1.5% of converters were backdoor in 2009, by 2015 it was more than 37%. Individuals who made conversions in 2010 could recognize the resulting income immediately or in equal amounts in 2011 and 2012. We estimate that approximately 31% of individuals made the ex post incorrect decision: they could have reduced their tax liability by making the opposite choice and the median extra tax liability was \$320. We find evidence in support of several Roth conversion motives: individuals are more likely to convert in a transitory low tax year, use Roth accounts as diversification against tax rate uncertainty, and, among individuals who contribute the maximum in other retirement accounts, use conversions to increase after-tax savings.

Keywords: Individual Retirement Arrangements, Roth IRA, Roth conversion, asset accumulation, income taxation, retirement

JEL Codes: D14, H24

*Goodman and Schramm: Joint Committee on Taxation, U.S. Congress. Corresponding author email: heidi.schramm@jct.gov. This research embodies work undertaken for the staff of the Joint Committee on Taxation, but as members of both parties and both houses of Congress comprise the Joint Committee on Taxation, this work should not be construed to represent the position of any member of the Committee.

1 Introduction

Traditional and Roth individual retirement arrangements (IRAs) are both tax-preferred savings accounts but they differ in how they provide tax benefits. Contributions to traditional accounts are generally tax deductible at the time of the contribution, earnings accrue with tax deferral, and the return of deductible contributions and earnings are taxed upon withdrawal in retirement.¹ In contrast, contributions to Roth accounts are after-tax, but both the earnings and return of contributions are tax free upon withdrawal in retirement. Prior to 2010, approximately 10% of individuals were unable to save in a Roth account because of the income limits associated with Roth contributions.² In addition, there were income limits associated with Roth conversions, which allow an individual to convert traditional assets to Roth, and pay tax on the taxable portion of the assets at the time of conversion. By 2009, 25% of individuals were unable to execute a Roth conversion and in 2010, the income limits associated with conversions were eliminated. We study Roth conversions using administrative tax data on all individuals who made conversions from 2003 through 2015. We document time trends in conversion behavior, investigate how the 2010 removal of the conversion income limits affected conversions in that year, show how the removal of these limits allowed for “backdoor Roth” contributions thereafter, and explore potential motives for making a Roth conversion.

The decision about whether to contribute to a traditional or Roth IRA, or do a Roth conversion once eligible, is complicated and depends on several factors. First, the decision depends on whether one anticipates being in a higher tax bracket in retirement relative to the present, in which case saving in the Roth IRA is more likely to be preferable. Second, an individual may desire to save in a Roth account in order to increase their after-tax consumption in retirement. The contribution limit is the same for both types of IRAs and applies to total contributions across all IRAs: for tax years 2015-2018 it is \$5,500 annually, with a limit of \$6,500 for individuals 50 or older.³ Therefore, an individual can save more, after tax, in a Roth IRA than in a traditional account. Third, because of considerable uncertainty about future income levels and federal tax policy, and thus relative tax rates,

¹There are income limits associated with the deductibility of traditional contributions which tend to make higher income individuals ineligible for this type of IRA contribution. However, there are no income limits for non-deductible traditional contributions.

²We calculate this figure using a nationally representative sample of individuals drawn from administrative tax data for years 2003 through 2009.

³If the individual’s taxable compensation for the year is below the contribution limit, their total contributions cannot exceed their taxable compensation.

there is a diversification motive to save in both types of accounts.⁴⁵ Fourth, individuals may be motivated to make a Roth contribution because they intend to make a penalty-free withdrawal of their basis prior to retirement: a similar withdrawal from a traditional account would incur a ten percent penalty prior to age $59\frac{1}{2}$. Finally, in contrast to traditional IRAs, Roth accounts are not subject to the required minimum distribution rules (RMDs), which require that a certain fraction of traditional assets are withdrawn each year, starting the year the account holder turns $70\frac{1}{2}$. For individuals who have a bequest motive for saving and would prefer not to draw down their IRAs, or prefer to withdraw them more slowly as compared to the required minimum distribution schedule, the Roth may be preferable.⁶ The decision of whether to convert traditional assets to a Roth depends on similar factors, as the optimal allocation of assets across pre- and post-tax accounts may change relative to when the contributions were initially made, for example due to updated information, expectations, and preferences, or becoming newly eligible to contribute to an after-tax account, as was the case for many high-income taxpayers starting in 2010.

With the passage of the Tax Increase Prevention and Reconciliation Act of 2005 (TIPRA; P.L. 109-222) on May 17, 2006, the income limits on Roth conversions were removed for tax years beginning in 2010 and later. The policy change allowed high-income individuals to convert traditional IRA assets which had accumulated over many years, which was the usual reason for doing a conversion prior to 2010. However, the policy change also enabled high-income individuals to perform a two-step transaction, known as a “backdoor Roth”, which approximates a direct contribution to a Roth IRA. To perform a backdoor Roth, an individual makes a non-deductible contribution to a traditional IRA of any amount up to the annual contribution limit. Shortly thereafter, the individual converts the entire amount to a Roth IRA. The individual recognizes income from the conversion only to the extent that the value of the traditional IRA appreciated between the time when the non-deductible

⁴Kotlikoff et al. (2008) study the relative benefits of saving in traditional and Roth accounts under various assumptions about future tax policy. They find in the absence of future tax changes, middle-income, single-parent households benefit slightly more from Roth accounts, while married and other single households benefit more from traditional accounts. However, their findings largely depend on the extent to which future tax rates change.

⁵A related motivation for after-tax saving is if individuals expect to have volatile income in retirement, in which case they may plan to withdraw from their Roth account in high-tax years and traditional in low-tax years to smooth taxable income, and thus total tax liability, over time.

⁶Mortenson et al. (2017) show that a large fraction of traditional account-holders prefer to withdraw less than their required minimum. They also find that among individuals age $70\frac{1}{2}$ and older, average and median distributions, measured as a fraction of prior year account balance, from Roth IRAs are substantially smaller than those from traditional IRAs.

traditional contribution was made and the Roth conversion.⁷ After the two-step process, the individual has roughly the same amount of assets in the Roth account as if she had made the original contribution to a Roth account. This strategy is allowed regardless of whether the taxpayer is income-eligible to make a direct Roth IRA contribution because there are no income limits for making non-deductible IRA contributions or the conversion.

We find a large response to the 2010 removal of the conversion income limits: while 200,000-250,000 individuals made Roth conversions annually between 2003 and 2009, just under 1 million made conversions in 2010. The aggregate value of conversions was between \$2 and \$3 billion annually prior to 2009, increased to \$6 billion in 2009, and was \$56 billion in 2010. In the years thereafter, about half a million individuals made conversions totaling between \$6 billion and \$15 billion annually. Starting in 2009, we estimate that an increasing share of converters are backdoor converters: while approximately 1.5% of converters were backdoor in 2009, by 2015 more than 37% of converters were backdoor.

Despite the sizable responses, only 3.3% of IRA-holders made a Roth conversion over the period of study: thus, while conversions are the focus of this paper, Roth conversions are the exception rather than the rule for IRA holders.⁸ We find Roth converters are more likely to be male, have higher savings (as proxied by total IRA assets, investment income, and workplace qualified plan saving), and have non-wage sources of income. Additionally, even prior to the conversion, Roth converters were more likely to have a Roth IRA, beyond what can be explained by differences in age, income, or marital status.

For Federal tax purposes, the income associated with a conversion is included in income in the year in which the conversion occurs. However, individuals who made Roth conversions in 2010 could elect to recognize the resulting income immediately in 2010, or in equal amounts in 2011 and 2012, which was the default option. The ability to split the recognition over two years might be attractive to taxpayers: holding income constant, the progressive rate structure means that total tax liability will generally be lower when income is recognized in a smoother manner over a course of years rather than unevenly over time. Furthermore, the opportunity to delay recognition might be attractive due to the time value of money. However, the optimal decision differed for every taxpayer based on their own preferences and beliefs about their marginal rates in 2010, 2011, and 2012.

We find that approximately 67% of individuals chose to delay recognition to 2011/2012. Using a simplified tax calculator, we estimate that approximately 31% of individuals made

⁷In practice, appreciation can be prevented by investing the IRA assets in cash. Furthermore, for this strategy to work as described, the taxpayer must not own any other IRAs with less-than-full basis. An individual that makes only deductible contributions to an IRA, or who rolls over amounts from an employer plan which was funded with pre-tax dollars, will have zero basis in any Roth conversion amount. As a result, the entirety of the Roth conversion amount will be taxable income in the year of the conversion. In contrast, non-deductible IRA contributions can cause the taxpayer to have positive basis in her IRA, reducing or eliminating income recognition upon the conversion.

⁸Specifically, among individuals who have an IRA for at least one year between 2003 and 2015, we observe 3.3% made at least one Roth conversion during the same period.

the ex post incorrect decision: they could have reduced their tax liability by making the opposite choice.⁹ Among individuals who made an ex post suboptimal decision, we estimate that the median extra tax liability was \$320.¹⁰ While we do not expect that ex ante optimal decisions based on realistic expectations about the future necessarily equate to ex post optimal decisions, we find this to be a useful exercise because of the relevance more generally for retirement savings decisions: provided an individual can make pre- and post-tax retirement savings contributions, the decision about which type to make critically depends on expectations about future tax rates relative to current rates. If individuals experience difficulty making optimal decisions over a two or three year horizon, it seems reasonable that they have an at least equally difficult planning problem over the course of the time between when they make contributions and when they draw down assets, which could span 30 or 40 years. There is no existing data that allows researchers to explore the optimality of the allocation of assets across pre- and post-tax accounts over such a long time horizon: we believe the findings presented here are valuable in thinking about an important issue that researchers cannot study directly.¹¹

We find evidence consistent with several explanations for why individuals do Roth conversions.¹² First, we consider whether a transitively low tax rate affects the probability of doing a conversion: if an individual's marginal tax rate is temporarily low in a given year, a Roth conversion is relatively more attractive even if there is little uncertainty regarding future tax rates. We find non-causal evidence of the relationship using a distributed lag framework, which shows the path of Roth conversions around marginal tax rate changes.¹³ We find that a transitory ten percentage point reduction in the marginal tax rate is associated with an increase in Roth conversions of about 25% of the baseline amount of conversions. This is consistent with taxpayers responding to marginal rate changes by adjusting their Roth conversion behavior. Additionally, the conversion probability returns close to baseline in the year after the change, suggesting that the increase in conversions is not a mere timing shift.

Second, we find evidence that backdoor converters are more likely to be constrained by

⁹This figure assumes a 6% discount rate. The analogous figure for a 0% discount rate is 29%.

¹⁰This figure assumes a 6% discount rate. The analogous figure for a 0% discount rate is \$500.

¹¹The 2010 income recognition timing decision is much simpler than decisions associated with the life-time allocation of retirement savings across pre- and post-tax accounts because the former only concerns expectations about relative marginal tax rates, while the latter involves all of those factors discussed in the second paragraph. Therefore, if a large fraction of individuals has difficulty with the simpler decision, it is reasonable to anticipate that at least as large a fraction will struggle to make an optimal decision in the more complicated case.

¹²We do not explore whether individuals are motivated to save in Roth IRAs to avoid RMDs, as Mortenson et al. (2017) establish sufficient evidence of this motive.

¹³A causal estimate requires exogenous variation in marginal tax rates, which is difficult to come by given the dependent variable of interest and the time period under consideration. While in 2013 there is tax rate variation for the top bracket, whose rate increased, there is not a satisfactory control group with whom to compare their conversion behavior. The rate increase was retroactive and, therefore, leading up to the change many taxpayers responded, including those who ultimately did not experience a rate increase.

contribution limits across different types of retirement savings accounts – employer plans and traditional IRAs. For example, we find that 61% of backdoor converters contribute the maximum allowable amount to their employer-provided retirement plan and the IRAs for which they are eligible, about 30 percentage points higher than other converters.¹⁴ This finding is consistent with individuals using backdoor Roth contributions to increase saving in tax-advantaged accounts when they are unable to do so through other channels.

Third, we consider whether individuals are motivated to do a conversion as a form of diversification: future tax rates are uncertain and some individuals may find it optimal to hold a mix of assets in traditional and Roth accounts. Indeed we find evidence that supports this motive: depending on the year, between 30% and 55% of Roth converters do not convert all of their traditional assets. Because individuals could smooth their conversions over several years in order to minimize the tax liability associated with a full account conversion, we use the panel nature of our data to test whether individuals slowly convert their assets over multiple years. We find that among individuals who did a conversion that was less than 100% of their traditional balance in 2010, by 2015 only about 15% had converted all of their assets to Roth, consistent with the hypothesis that some individuals are motivated to save in Roth IRAs as a form of diversification to protect against future tax rate uncertainty.¹⁵

Fourth, we examine the extent to which individuals are motivated to make a conversion because they plan to withdraw some portion of the account prior to retirement. Specifically, an individual who performs a Roth conversion may make a penalty-free withdrawal from the Roth IRA up to the amount of the conversion after a 5-year waiting period has passed, even if the withdrawal occurs prior to age $59\frac{1}{2}$; a similar withdrawal from a traditional IRA would incur a ten percent penalty (in addition to being subject to tax as usual). Recent research shows a sizable response to the 10% penalty: Goda et al. (2018) find that among individuals of ages around the $59\frac{1}{2}$ age threshold, there is an 80% increase in annual withdrawals upon reaching the age cutoff.¹⁶ However, we find little evidence that Roth IRA distributions increase five years after the date of conversion: it does not seem to be the case that penalty-free withdrawals were a major motive behind Roth conversions.¹⁷

Lastly, we explore three pieces of evidence which suggest that information plays a role in conversion behavior. First, despite the fact that the policy change did not take effect

¹⁴Other converters are individuals who execute taxable conversions – i.e. conversions that result in taxable income. This figure includes individuals who are not able to contribute to an employer-provided plan because their employer does not offer one. Among individuals who have access to such plans, 62.5% of backdoor converters make contributions at the maximum compared to 25% among taxable converters.

¹⁵It is also possible that a large number of individuals plan to slowly convert all traditional assets over more than 6 years, in which case we would underestimate the total fraction of full converters and overstate the case for the diversification motive.

¹⁶Additional research shows that leakage from retirement accounts in general is substantial: approximately \$0.40 of every \$1 contributed into an account prior to age 55 is withdrawn from the account prior to retirement (Bryant et al., 2010; Argento et al., 2015).

¹⁷We also do not find an increase after 6 years.

until 2010, Roth conversions increased substantially in 2009 relative to several years prior. This is consistent with taxpayers responding in 2009 to news stories about the impending policy change, though this fact can be explained by other factors as well.¹⁸ Second, we find that there was a large increase in the number of Roth converters in 2010 even among the set of taxpayers with low enough income (i.e., below the \$100,000 conversion AGI threshold) to be unaffected by policy. Specifically, the number of Roth converters with income below \$90,000 increased from 230,000 in 2009 to 410,000 in 2010, a substantial increase over prior years. This is further evidence that taxpayers may have responded to news or other forms of information provision coincident with the policy change. Finally, we exploit discontinuities in exact date-of-birth and find that being subject to RMDs reduces the likelihood of making a Roth conversion by 20%. While individuals cannot make traditional IRA contributions after age $70\frac{1}{2}$, there is no prohibition on making a Roth conversion after that age.¹⁹ ²⁰ This result can most easily be explained by assuming that many older taxpayers find the rules regarding IRA contributions, Roth conversions, and RMDs – and the interactions thereof – to be confusing.

There is a lack of research that analyzes the effects of the 2010 removal of the Roth conversion income limits. This is likely due to data limitations: Roth conversions are still relatively rare in the general U.S. population, with just under 1 million individuals making a conversion at the peak of activity in 2010. Household surveys do not ask about conversion behavior and, therefore, administrative data from tax records or plan providers is required to get a thorough picture of behavior.²¹ Bryant and Gober (2013) use data from the Internal Revenue Service (IRS) Statistics of Income to show the aggregate increase in conversions from 2004 through 2010 and the fraction of taxpayers who made a conversion in 2009 and 2010 by AGI categories.²² Our data provide us with a unique opportunity to study a previously undocumented and increasingly important type of retirement savings decision – backdoor Roth contributions. To our knowledge, we are the first researchers to document the prevalence of backdoor Roth contributions, estimate the ex post optimality of the income recognition timing decision associated with 2010 conversions, and empirically explore several reasons that individuals may prefer saving in Roth IRAs using a nationally representative source of data.

¹⁸The most likely alternative explanation is that conversions increased in 2009 because it was a low marginal rate year for many taxpayers.

¹⁹There is no age limit associated with Roth contributions. Individuals can also make rollover contributions to traditional and Roth IRAs regardless of age. The IRS Publications 590-A and 590-B contain more details about IRA rules.

²⁰Because individuals cannot make traditional IRA contributions after age $70\frac{1}{2}$, only individuals who have an existing traditional IRA, or create a new one because of a rollover from a defined contribution account, are able to make a conversion after the age. Therefore, the sample of individuals included in this analysis is limited to those who have a traditional account in the prior year.

²¹Hoffman (2007) and Burman (2006) thoroughly discuss the potential savings behavior and tax revenue effects of the 2010 policy change, respectively.

²²Their findings for this period are similar to our own.

In related research, Beshears et al. (2014) and Beshears et al. (2017) study Roth contribution behavior following the introduction of Roth options in employer-sponsored retirement plans.²³ Beshears et al. (2014) use retirement plan data from twelve companies that introduced the option between 2006 and 2010 and find the following: a year after the Roth plan introduction, about 9% of plan participants have positive Roth balances, Roth contributions constitute 5.4% of retirement plan contributions, and 55% of employees who contribute to a Roth also contribute to a traditional account, consistent with a tax diversification motive. Beshears et al. (2017) use similar plan data from eleven companies and find that contribution rates do not differ between employees hired after a Roth option was introduced and those hired before. This finding, which is inconsistent with a simple neo-classical model of lifetime savings, implies that the amount of retirement consumption purchased by 401(k) contributions increases after the introduction of the Roth option.

While the Beshears et al. (2014; 2017) studies are related to the research presented here in that we both study responses to a new Roth option, our findings differ in several important ways. First, our research is primarily concerned with analyzing the 2010 conversion policy change and the backdoor Roth contributions that have continued since. Defined contribution (DC) plans with a Roth option are different from Roth IRAs in several ways: DC plans do not have income limits associated with contributions, have much higher contribution limits, and are subject to the RMD rules.²⁴ While they are both after-tax retirement savings accounts, Roth IRA and DC Roth accounts are distinct and we do not necessarily expect the individuals who use them or the way they are used to be the same. Second, the Roth conversion policy change was at the federal level, thus affecting all taxpayers, whereas not all DC plans offer a Roth option.²⁵ Employees with access to plans that offer a Roth option are not representative of all employees who have access to DC plans: because employees who desire a Roth option are most likely to pressure their employers to offer it, the employees of those plans that do not yet offer it may have a different allocation of assets across pre- and post-tax accounts upon introduction of a Roth option relative to those who already have it.²⁶ In contrast, we use data from all individuals who made a Roth conversion between 2003 and 2015, supplementing with a nationally representative sample of all taxpayers as necessary. Finally, because the 2010 change was at the federal level, it garnered substantial financial

²³The Economic Growth and Tax Relief Reconciliation Act of 2001 allowed employer retirement plan sponsors to add a Roth 401(k) option to defined contribution plans beginning in tax year 2006.

²⁴They differ in other ways that tend to depend on the specific plan that offers the 401(k). For example, the employer match, which is pre-tax, may affect 401(k) saving rates. In addition, the rules associated with borrowing from the accounts differ.

²⁵According to the Plan Sponsor Council of America (2012), 49% of 401(k) plans offered a Roth option in 2011. This figure has likely increased since then, as Vanguard reports that in 2016, 60% of their defined contribution plans offered the option.

²⁶Beshears et al. (2017) show that the eleven companies for which they use data are large, ranging from 10,000 to 100,000 employees each, seven of them are in the financial services industry, and four had average salaries that exceeded \$100,000 in 2010. The authors note that employees at the firms are likely more financially sophisticated than the typical U.S. employee.

news coverage which we find may have affected who responded, as previously discussed.

The remainder of the paper proceeds as follows. In Section 2, we discuss the 2010 policy change in detail. Section 3 describes our tax data. Section 4 establishes several stylized facts about Roth conversions: time trends in conversion behavior from 2003 through 2015, backdoor Roth contributions beginning in 2010, and the income recognition decisions associated with 2010 conversions. In Section 5, we examine various characteristics of Roth converters over time. In Section 6, we present evidence consistent with several explanations for why individuals do Roth conversions. The last section concludes.

2 Policy Background

There are two types of individual retirement arrangements (“IRAs”): traditional IRAs and Roth IRAs.²⁷ Contributions to a traditional IRA may or may not be deductible from taxable income in the year they are made, depending on the taxpayer’s modified adjusted gross income (MAGI) and whether the individual (or the individual’s spouse) is an active participant in an employer-sponsored retirement plan. Distributions from a traditional IRA are includible in gross income to the extent they are not attributable to a return of non-deductible contributions. Roth IRAs were established as another type of retirement savings account in 1997 with the enactment of the Taxpayer Relief Act of 1997. For a Roth IRA, an individual makes after-tax contributions and does not pay tax on qualified withdrawals.²⁸

There are income limits associated with making deductible traditional IRA and Roth contributions. For example, in tax year 2018, individuals who are covered by a retirement plan at work and file as married filing jointly with a MAGI of \$101,000 or less can make a full deduction of their traditional IRA contributions, while those with a MAGI between \$101,000 and \$121,000 get a partial deduction, and those with more than \$121,000 can only make nondeductible contributions.²⁹ In tax year 2018, individuals who file as married filing jointly with a MAGI of less than \$189,000 can make a Roth contribution, with a phaseout ending at \$199,000: individuals with a MAGI of greater than or equal to \$199,000 cannot make Roth contributions.

These contribution limit patterns are adjusted for inflation but have been broadly similar over time. Therefore, prior to 2010 when the conversion limits were removed, high income households with access to employer-sponsored retirement plans could not save in Roth IRAs and were limited to making non-deductible contributions to traditional IRAs. The only

²⁷There are also SEP and Savings Incentive Match Plan for Employees (SIMPLE) IRA plans that are primarily used by self-employed and small business employees. The rules associated with SEP and SIMPLE IRAs are slightly different from those discussed here. We do not analyze contributions to these plans.

²⁸See IRS Publication 590-B for additional details about distributions from IRAs.

²⁹Different income limits apply depending on filing status and whether an individual, and their spouse if married, are covered by an employer-provided retirement plan. See IRS Publication 590-A for additional details about contributions to IRAs.

tax advantage associated with non-deductible contributions to traditional IRAs is the tax deferral on earnings that accrue within the account. The earnings in such an account are taxed at ordinary rates upon withdrawal in contrast with a Roth, where the earnings are tax-free along with the return of after-tax contributions. By our estimation, approximately 12% of individuals were ineligible to make Roth contributions in 2015.

Prior to 2010, there were also income limits associated with converting assets saved in a traditional IRA to a Roth. The conversion limits were lower than the Roth contribution limits. For example in tax year 2009, individuals with a MAGI of more than \$100,000 were not able to make conversions. In that year, direct Roth contributions completely phased out at a MAGI of \$176,000 for taxpayers filing as married filing jointly and \$120,000 for single and head of household filers. We estimate that in 2009, 25% of individuals were ineligible to make a Roth conversion. With the passage of TIPRA on May 17, 2006, the income limits on Roth conversions were removed for tax years beginning in 2010.

The amount converted is treated as a distribution from the traditional IRA for income tax purposes, except that the 10-percent additional tax on early withdrawals does not apply. For individuals who make deductible contributions or have positive earnings after making non-deductible contributions to a traditional IRA, the conversion will result in taxable income. If an individual has more than one IRA and any of their IRAs have tax basis (*i.e.* non-deductible contributions were made), the determination of the income resulting from the conversion is made by treating all traditional IRAs as one account and all distributions as one distribution. This means that individuals cannot selectively convert accounts that have high basis. Therefore, for individuals who contributed or rolled over from employer plans substantial pre-tax amounts in their traditional IRA(s), the taxable income associated with a conversion could be quite large. Accounting for basis across all traditional IRAs could serve as a deterrent to executing a conversion.

For Federal tax purposes, the income associated with a conversion is generally includible in income in the tax year in which the conversion occurs. However, for conversions that occurred in 2010, the default treatment of the income associated with the conversion was that it would be recognized in equal amounts in tax years 2011 and 2012. Individuals were also permitted to recognize the resulting income immediately in 2010. The individual tax rate reductions under the Economic Growth and Tax Relief Reconciliation Act of 2001 (EGTRRA) and the Jobs and Growth Tax Relief Reconciliation Act of 2003 (JGTRRA) were originally set to expire at the end of 2010, which may have made immediate recognition preferable. However, the rate changes were temporarily extended in 2010 through the Tax Relief, Unemployment Insurance Reauthorization, and Job Creation Act of 2010, which was signed into law on December 17, 2010. For tax year 2010, individuals had until filing in April 2011 to decide if they wanted to pay tax on the amounts associated with 2010 conversions in 2010. Because individuals had several months to observe that tax rates did not increase after 2010, we do not expect that the potential increase in tax rates for years after 2010, as

observed during 2010, had an effect on the timing decision.³⁰

The rate cuts were made permanent, except for the top tax bracket, under the American Taxpayer Relief Act of 2012, which was passed on January 2, 2013 and was retroactive to January 1, 2013. Therefore, holding income constant, only the highest income households (single filers with AGI more than \$400,000 and married filing jointly filers with AGI more than \$450,000) were exposed to tax rate increases during the period of study, and those tax rates went into effect beginning in tax year 2013. We discuss how this may have affected conversion behavior in 2012, including among taxpayers who ultimately did not experience a rate increase, in Section 6.

3 Data

We use administrative tax data maintained by the IRS.³¹ The underlying database contains the near-universe of certain tax forms (such as Form 1040) and information returns (such as Form W2) over the period which we study. We observe detailed information about transfers to and from IRAs, linked to detailed information about each individual holder.

Our primary dataset is a panel dataset of two types of IRA holders. First, we include the full universe of individuals that ever make a Roth conversion (as indicated by the trustee of the recipient IRA on Form 5498) between 2003 and 2015.³² This allows us to characterize the entire distribution of individuals who make Roth conversions, including individuals with Roth conversion amounts in the long upper tail of the distribution of conversions who are unlikely to be captured in samples. Second, we include a 2.5% sample of all individuals who ever hold an IRA (as indicated by receipt of Form 5498) between 2003 and 2015. In most of the empirical analysis below, we apply sample weights such that the sample is representative of this “ever-IRA” population from which the 2.5% sample is drawn.³³

For these individuals, we create a nearly-balanced 13-year panel from 2003 to 2015, hereafter referred to as the “base sample.” Individuals enter the sample at age 18, and exit the sample at death or age 106.³⁴ We attach information derived from the following information returns: Form W2 (wages and contributions to qualified retirement plans, such as 401(k)s), Form 5498 (IRA contributions and balances), Form 1099-R (distributions from retirement plans, including IRAs), and Form 1099-SSA (Social Security benefits). Individuals

³⁰This may be particularly true for well-informed taxpayers: individuals who would have based their decision on the potential expiration in 2010 would also have been especially likely to observe that the rate cuts were extended.

³¹These data are pre-audit: they reflect reported information before any IRS audits.

³²This includes the set of individuals who subsequently recharacterize all or a portion of the conversion. However, in all of our results, our definition of “Roth conversion” is the conversion amount net of recharacterizations.

³³Mechanically, the Roth converter population is a subset of the ever-IRA population.

³⁴Because death is not perfectly reported to the Social Security Administration, it is reasonable to approximate that individuals over the age of 105 are deceased.

may receive several information returns of the same type in the same year. For example, individuals receive one Form 5498 for each IRA that they own in a given year. Thus, we extract the data from each information return and take the sum across the relevant fields for a given individual and tax year.³⁵ We include the information returns of the individual's spouse as separate variables if he or she is married.

We also attach information derived from Form 1040 and its associated schedules, which provides information on the level and sources of income for individuals in our sample. Additionally, we use information from Form 8606, which is required to be filed by individuals who make a Roth conversion, who make a non-deductible IRA contribution, or who fit a number of other circumstances. We supplement these data with Social Security Administration data on gender and exact dates of birth and death.

For the sake of descriptive statistics, we create a secondary 0.5 percent sample that is representative of the tax filing population of U.S. adults between 2003 and 2015. We include only individual-year observations in which the individual appears on *any* tax form, including information returns and tax returns for primary filers, secondary filers, and dependents. As shown by Cilke (2014), this restriction best approximates the U.S. resident population.³⁶

Table 1 presents summary statistics for each of the three samples. The IRA sample tends to be about 6 years older than the general population, with an average age of about 53 years old. The sample of converters has a similar mean age to the population as a whole, but with less variance, as a smaller share of the converter sample is age 65 or older. Unsurprisingly, the IRA sample and especially the Roth converter sample are positively selected. They tend to have higher income, as indicated by the higher filer share, the lower share with income eligible for an IRA deduction, and the higher 25th, 50th, and 75th percentiles of adjusted gross income.³⁷ Similarly, the IRA sample and the converter sample are substantially more likely to be married than the general population. By construction, the IRA sample and the Roth converter sample are much more likely to own an IRA (Roth or traditional) and to have made a Roth conversion.

Figure 1 shows the share of the population that was ineligible to make deductible IRA contributions, Roth IRA contributions, and Roth conversions from 2003-2015.³⁸ The share of individuals who were ineligible to make deductible IRA contributions is between 21% and 28% over the period, with a larger share of ineligible individuals in the early years of our sample. The share ineligible to make direct Roth contributions is relatively low, around 10% on average, and slightly increasing since 2009. The share ineligible for Roth conversions

³⁵We are careful to drop duplicate and/or corrected information returns to the best of our ability.

³⁶In order to manage the size of the dataset, this population dataset contains only a subset of the base sample variables, but all of the base sample variables could in principle be used in this population sample.

³⁷Note that, for the purpose of this table, we code AGI to be equal to zero for nonfilers.

³⁸Because we are primarily interested in the effect of removing the conversion income limits, the definition of ineligible used here is with respect to the MAGI limits, and their interaction with filing status and an individual's own or their spouse's access to employer-provided retirement plans. It does not take into account taxable compensation requirement associated with contributions.

was increasing from 2003 to 2009, at which point it was approximately 25%. Prior to 2010, the share ineligible for a Roth conversion is larger than the share ineligible for a direct Roth contribution due to the fact that the AGI limits for conversions were lower than for contributions.

Only 3.3% of the IRA sample has ever made a Roth conversion: thus, while it is the focus of this paper, Roth conversions are the exception rather than the rule for IRA holders. Despite removal of the income limits, the take-up rate of Roth conversions was low among newly eligible individuals following the 2010 policy change. Conditional on having a traditional or Roth IRA in the prior year, only 0.3% of individuals made a conversion prior to 2010, 1.7% made a conversion in 2010, and between 0.7-0.9% made a conversion annually from 2011 to 2015.³⁹ Among the group of individuals most likely to make a conversion – those whose AGI is too high for a direct Roth contribution – the share who made a backdoor Roth contribution after 2010 was comparatively much higher: 0.7% made a backdoor Roth contribution in 2010 and by 2015, 1.6% of individuals made a backdoor Roth contribution. In aggregate, Roth converters represent a small fraction of taxpayers in any year.

3.1 Calculation of Basis & Backdoor Conversions

We are interested in the basis that taxpayers have in their IRAs because it is a key factor in individuals' decisions to convert some or all of their traditional IRA into a Roth IRA, as discussed in Section 2. However, while basis in the converted assets is reported on Form 8606 among those who convert, we are interested in the IRA basis of non-converters as well.⁴⁰ We calculate basis at the end of the year (b_t) as follows. First, we assume that basis is zero at the end of 2002 (or the beginning of 2003); i.e., $b_{2002} = 0$. Thereafter, we determine basis recursively, according to the following equation, where V_t denotes the fair market value of all (traditional, SEP, and SIMPLE) IRA assets at the end of time t , c_t denotes non-deductible IRA contributions, d_t denotes IRA distributions (not including rollovers or Roth conversions), and x_t denotes Roth conversions:⁴¹

³⁹Because these figures pertain to individuals who had any type of IRA in the prior year, the population averages are much smaller.

⁴⁰Additionally, it turns out that many converters fail to file Form 8606.

⁴¹In most cases, we calculate non-deductible contributions as contributions reported on Form 5498 minus the IRA deduction claimed on Form 1040. In the unusual case when both spouses of a jointly-filing couple each make an IRA contribution, with an IRA deduction between 0 and the sum of the contributions, we use a more complicated procedure. If the IRA deduction equals the contribution of one spouse (and the contributions of both spouses are unequal), we assume that the other spouse's contribution was entirely non-deductible. If the IRA deduction exactly equals the amount allowable under the statutory application of income phase-outs of one or both spouses, we assume that the residual is non-deductible. Finally, in other cases, we pro-rate the non-deductible share of the couple's IRA contributions in proportion to their total IRA contributions.

$$\begin{aligned}
b_t &= b_{t-1} + c_t - \lambda_t(d_t + x_t) \\
\lambda_t &= \frac{b_{t-1} + c_t}{V_t + d_t + x_t}
\end{aligned}
\tag{1}$$

λ_t represents the “basis share” of total IRA assets (which, for this purpose, adds back the distributions and conversions made during the year). Thus, basis at the end of time t is equal to basis at the end of time $t - 1$, plus nondeductible contributions, minus the basis share of distributions. This rules out complications such as (1) non-deductible contributions made for time t between January 1 and April 15 of time $t + 1$, (2) corrective distributions that affect basis, and (3) the special calculation that must be made for individuals who take a distribution and make a nondeductible contribution in the same year. Given λ_t , we calculate the basis in the amount converted as equal to $\lambda_t \times x_t$. Finally, we define a “backdoor conversion” to be a conversion where $\lambda_t \geq 0.8$.

In Appendix Figure A1, we verify that this procedure works reasonably well among the subset of converters that file Form 8606 in the year of conversion, as required. In the top panel, we plot the aggregate dollars of basis in Roth conversions reported by taxpayers over time (solid line), as well as the aggregate dollars of basis in this subsample as calculated by our procedure (dotted line). The qualitative pattern in both series is similar: basis is very small prior to 2010, spikes in 2010, and then falls to a higher baseline level afterward. However, we calculate somewhat less basis (especially in 2010) relative to taxpayers’ direct reports. This discrepancy could arise because taxpayers are making mistakes (e.g., failing to apply aggregation rules which aggregate all traditional IRA assets into one bucket) or because our procedure does not capture all basis (e.g., because basis was not zero prior to 2003). In the bottom panel, we show a similar plot. We plot the number of backdoor converters over time in this subsample, using basis reported by taxpayers (solid line) and calculated by us (dotted line). As in the top panel, the two series have a similar path, though taxpayer-reported basis tends to find slightly more backdoor converters than we find using the basis calculated by us. Nevertheless, we interpret this figure as showing that our basis-calculation procedure is picking up the most important patterns. To the extent we underestimate basis, we will undercount the fraction of converters that are backdoor converters.

3.2 Calculating tax liability under alternate scenarios

We use a simplified tax calculator to calculate tax liability of individuals under various scenarios. We use this calculator for two purposes. First, we use this calculator to determine the tax liability owed by 2010 converters on the conversion amount if (1) the amount is recognized immediately in 2010 or (2) the amount is recognized half in 2011 and 2012.⁴²

⁴²See Section 4.3 for further details.

Second, we use this calculator to determine marginal tax rates on the first dollar of IRA distributions.⁴³ This calculator takes into account the following features of the federal individual income tax as it existed during the 2003-2015 period: the ordinary income tax rates and brackets (including personal exemptions, and standard deduction/itemized deductions), the interaction between ordinary income and the statutory tax rate on capital gains, the partial taxation of Social Security benefits, and the alternative minimum tax.⁴⁴ This calculator abstracts away from such features of the tax code as the overall limitation on itemized deductions, the phaseout of personal exemptions, and other itemized deductions (e.g., the medical expense deduction) or credits (e.g., the Child Tax Credit and the Earned Income Tax Credit) which are a function of income.

4 Stylized Facts About Roth Conversions

4.1 Roth Conversions from 2003-2015

In this section, we establish that Roth conversions increased substantially in 2010 relative to prior years. After 2010, aggregate Roth conversions fell to a higher baseline level, at least twice as much as pre-2010 levels annually, on average. Figure 2 shows the number of conversions and the total dollar amount of these conversions in 2009 dollars.⁴⁵ The figure shows that roughly 200,000 - 250,000 individuals made Roth conversions in each year between 2003 and 2009. This jumps to 950,000 individuals in 2010, and falls to 400,000-550,000 individuals in each of the years between 2011 and 2015. The series of dollar value of conversions shows a similar pattern. Between 2003 and 2008, aggregate Roth conversions were between \$2 billion and \$3 billion annually. The aggregate value of conversions increased to \$6 billion in 2009 and \$56 billion in 2010. After 2010, the value decreased to \$9 billion in 2011, increased to \$15 billion in 2012, and decreased to between \$6 billion and \$9 billion from 2013 to 2015. The higher number of converters and dollar amounts of conversions after 2010, relative to pre-2010 numbers, are likely due to backdoor Roth conversions, as discussed in the next section.

We believe that the increase in conversions, in particular the dollar amounts, between 2011 and 2012 was due to anticipation of the expiration of EGTRRA and JGTRRA after 2012. This hypothesis is consistent with Saez (2017), who finds other behavioral responses associated with the 2013 tax increase, which raised top marginal rates on capital income by 9.5 percentage points and on labor income by 6.5 percentage points. Saez (2017) estimates a short-run elasticity of reported income with respect to the net-of-tax rate larger than

⁴³See Section 6.1 for further details.

⁴⁴For purposes of calculating the alternative minimum tax, we assume that the difference between the taxpayer's initial AGI and initial AMT income is unchanged by any distribution.

⁴⁵Conversions are shown on Form 5498 and are adjusted by subtracting recharacterizations of Roth contributions, as indicated on Form 1099-R.

one, due to income retiming for tax avoidance purposes. In Section 6, we discuss evidence of a strong (non-causal) relationship between low transitory marginal tax rates and Roth conversions.

Aggregate conversions also increased substantially in 2009, which may have been a low marginal rate year for many individuals: the business cycle could have caused income to decrease, pushing taxpayers toward the lower end of the progressive rate schedule. However, as discussed in Section 6, this change is also consistent with an informational response to financial news stories that covered the 2010 policy change in advance of the elimination of the income limit. Because we find conversions increased in 2010 among individuals who could have made direct Roth contributions prior to 2010, we believe information may be responsible for part of the increase in conversions in 2010 and may also account for the increase in aggregate conversions in 2009.

Figure 3 plots the mean and median conversion amount over time. Throughout the sample period, the mean was substantially larger than the median, reflecting the skewed nature of Roth conversions. In most years, this ratio is about three, though it increases to more than four in 2012, perhaps reflecting the response of high-income taxpayers to the potential expiration of EGTRRA and JGTRRA. Both the mean and median conversion amounts roughly double in 2010 relative to 2009.

4.2 Backdoor Roth Contributions

While there are income limits on direct Roth contributions, there are no income limits on Roth conversions since 2010.⁴⁶ Therefore, one common strategy for high income households who would otherwise not be able to contribute to a Roth is known as the “backdoor Roth.” To perform a backdoor Roth, an individual makes a non-deductible contribution to a traditional IRA of any amount up to the annual maximum.⁴⁷ Shortly thereafter, the individual converts the entire amount into a Roth IRA. In addition to the tax associated with making an after-tax contribution, the individual will recognize additional income only to the extent that the value of the IRA appreciated between the time of the non-deductible traditional contribution and the Roth conversion.⁴⁸ After the two-step process, the individual has roughly the same amount of assets in the Roth account as if she had made the original contribution to a Roth account. This strategy is allowed regardless of whether the taxpayer is income-eligible to make a Roth IRA contribution because there are no income limitations for the conversion or for making non-deductible IRA contributions.⁴⁹

⁴⁶In tax year 2018, the Roth contribution AGI phase-out for married filing jointly households is \$189,000 to \$199,000. For single and head of household filers, the AGI phase-out is \$120,000 to \$135,000.

⁴⁷In tax year 2018, the annual maximum is \$5,500 with a \$1,000 catch-up limit.

⁴⁸In practice, appreciation can be prevented by investing the IRA assets in cash.

⁴⁹However, an individual generally may not make a an IRA contribution, deductible or non-deductible, in excess of her taxable compensation for the year, with special rules for joint filers.

For this strategy to work as described, the taxpayer must not own any other IRAs with less-than-full basis. An individual that makes only deductible contributions to an IRA, or who rolls over amounts from an employer plan which was funded with pre-tax dollars, will have zero basis in any Roth conversion amount. As a result, the entirety of the Roth conversion amount will be taxable income in the year of the conversion. For individuals who have been making pre-tax retirement savings contributions for a large fraction of their careers, the includible income associated with a Roth conversion could be quite large and represent a significant barrier to converting. In contrast, non-deductible IRA contributions can cause the taxpayer to have positive basis in her IRA, reducing or eliminating income recognition upon the conversion.

Once the income limits on Roth conversions were lifted in 2010, the backdoor Roth contribution strategy may have been an attractive option for taxpayers who desired to make Roth IRA contributions but were previously income-ineligible. As discussed in Section 3, we estimate that about 10% of individuals are ineligible to make direct Roth contributions in any year due purely to the income limits. In addition, because the policy change was made law in May 2006, it might have been attractive to make non-deductible contributions over the course of several years between 2006 and 2010. While the longer time period would have exposed taxpayers to “appreciation risk,” as they would be required to recognize income on the earnings of the IRA between contribution and conversion, the earnings associated with non-deductible contributions between 2006 and 2010 was probably modest for most taxpayers.⁵⁰ We do not observe an increase in non-deductible contributions starting in 2006.⁵¹

Figure 4 describes the aggregate amount of Roth conversion basis. The top panel plots aggregate basis, in dollars, on the same graph as aggregate Roth conversion dollars.⁵² The bottom panel shows the ratio of the two series – i.e., the basis share of total conversions. The bottom panel shows that the basis share of conversions in the years immediately prior to 2010 was between 5% and 7%. This increases to about 9% in 2010, and continues to increase thereafter – with a decrease in 2012 – reaching nearly 20% by 2015. This is certainly consistent with the some taxpayers engaging in the backdoor Roth strategy, as after-tax contributions that are immediately converted will have close to 100% basis if there are no additional pre-tax assets. However, as the top panel makes clear, basis represented very little of the 2010 spike in total conversions when measured in dollars. Low basis in a large traditional account would result in a large increase in taxable income upon a Roth conversion:

⁵⁰As discussed by Hoffman (2007), assuming an individual made annual, non-deductible contributions up to the limit for tax years 2006 through 2010 and received an 8% annual return, the income associated with a conversion in 2010 would be \$2,843 for an individual under 50 years of age, or up to \$3,492 for a taxpayer older than 50 in every year starting in 2006. The amount includible in income could be substantially larger if an individual held other IRAs with low or no tax basis when making a conversion.

⁵¹Aggregate data on non-deductible contributions is available by request.

⁵²See Section 3 for a discussion of how we calculate basis.

it is clear that among those who did make conversions, basis was low. Therefore, low basis was not a deterrent to converting for those who converted.

Figure 5 tells a qualitatively similar, but quantitatively distinct story. The top panel plots the number of backdoor converters on the same scale as the total number of converters.⁵³ The bottom panel plots the ratio of these two series. Figure 5 confirms that the vast majority of the spike in 2010 converters was not driven by backdoor converters. Starting in 2009, an increasing share of converters are estimated to be backdoor converters: while approximately 1.5% of converters are backdoor in 2009, by 2015 it is more than 37%. This is roughly double the share of total conversions represented by basis, reflecting the fact that backdoor conversions tend to be smaller than other conversions. This is reasonable given that the annual contribution limit ranges from \$3,000 to \$5,500 over the period.

4.3 Income Recognition Associated With 2010 Conversions

Individuals who made Roth conversions in 2010 were permitted to elect to recognize the resulting income immediately in 2010, or in equal amounts in 2011 and 2012. The possibility to split the recognition into two years might be attractive to many taxpayers: holding total income constant, the progressive rate structure means that total tax liability will generally be lower when income is recognized in a smoother manner over a course of years, rather than unevenly over time. Furthermore, the opportunity to delay recognition might be attractive due to the time value of money. However, this is a decision that will differ for every taxpayer based on their own preferences and beliefs about their marginal rates in 2010, 2011, and 2012.⁵⁴

In this section, we analyze the decision of when to recognize the income associated with a 2010 conversion. In order to determine whether recognition was immediate or delayed, we make several data restrictions. First, we restrict attention to the approximately 65% of 2010 converters who correctly filed Form 8606 in 2010, as this is the form where the recognition decision is directly observed. Second, we drop those who report having zero taxable amount (i.e., whose conversion amount is covered entirely by basis), as the timing decision is not relevant for this group. Third, we drop converters who report on their 2010 Form 8606 taxable amounts in 2010, 2011, and 2012, that are not valid options: specifically, we require (1) the total reported taxable amounts in 2010, 2011, and 2012 to approximately equal the total taxable amount (as recorded on Line 18 of that form) and (2) the entire taxable amount to be recorded either in 2010 or in approximately equal installments in 2011

⁵³As discussed in Section 3, we define a backdoor converter to be a converter whose basis share of the conversion amount is at least 80%.

⁵⁴The decision could also depend on liquidity constraints but we abstract from such a discussion here.

and 2012.⁵⁵ This leaves us with 382,000 individuals out of the 953,100 individuals that made a Roth conversion according to Form 5498.⁵⁶

In order to determine the optimality of the income recognition timing decision, we apply our simplified tax calculator (discussed in Section 3) to calculate tax owed in 2010, 2011, and 2012, under each of the two scenarios: immediate recognition or delay until 2011 and 2012. For example, if an individual recognized income associated with the conversion in 2010, we calculate what the tax liability associated with the conversion was and we calculate what their liability would have been had they instead spread the income recognition across 2011 and 2012. We define optimal as timing the income recognition such that the tax associated with the conversion is minimized. For this analysis, we use the taxable amount of the Roth conversion as reported by the taxpayer on Form 8606.⁵⁷ Given that we can only conduct this analysis for about 40% of the population of 2010 Roth converters, care should be taken in extrapolating the findings to the entire population of converters.⁵⁸

It is important to keep in mind that individuals chose which year to recognize the income associated with the conversion in 2010, or at the latest by filing in April of 2011. There is uncertainty in predicting what your taxable income, tax bracket, and even what tax rates will be two years later. The policy uncertainty during this period was particularly important because of the potential expiration of the Bush tax cuts at the end of 2012, as discussed in Section 2. Therefore, we do not expect that ex ante optimal decisions based on realistic expectations about the future necessarily equate to ex post optimal decisions. However, we find this to be a useful exercise because of the relevance more generally for retirement savings decisions: provided an individual can make pre- and post-tax retirement savings contributions, the decision about which type to make critically depends on expectations about future tax rates relative to current rates. If individuals experience difficulty in making an optimal decision over a two or three year horizon, it seems reasonable that they have an at least equally difficult planning problem over the course of the time between when they make

⁵⁵For this analysis, due to data availability, we must also drop the very small number of such taxpayers who file a paper return. Additionally, the database records whether the taxpayer checked the box indicating that he or she is recognizing the conversion income immediately. We have found this check box to be less reliable than the monetary amounts.

⁵⁶Of the 953,100 individuals who make a Roth conversion in 2010 (according to Form 5498), we estimate that 625,300 filed Form 8606 and reported a positive Roth conversion amount. We find 617,400 of these individuals in the database of electronic returns. Of these, 489,300 individuals reported a positive taxable amount. Of these, 382,000 recorded taxable amounts that made logical sense. Relative to the full set of converters with positive taxable amount, the subset used in this analysis is about one year younger, 1.3 percentage points more likely to be married, has median income that is \$9,000 larger, and median conversion amount \$2,000 larger (\$22,000 relative to \$20,000).

⁵⁷While the taxpayer may have calculated this amount in error, for example due to an incorrect calculation of basis, we assume the taxpayer will generally make the decision regarding the timing of recognition based on this reported amount.

⁵⁸The fact that nearly 35% of 2010 converters did not correctly file Form 8606 indicates considerable confusion associated with the policy change, or IRA policy in general. As discussed in Section 6, there is evidence of confusion with respect to other aspects of Roth conversions.

contributions and when they draw down assets, which could span 30 or 40 years. There is no existing data that allows researchers to explore the optimality of the allocation of assets across pre- and post-tax accounts over such a long time horizon: we believe the findings presented here are valuable in thinking about an important issue that researchers cannot study directly.⁵⁹

In Table 2, we show the cross-tabulation of the decision by the taxpayer to recognize immediately or delay (across the columns) against the ex-post optimality of doing so (across the rows). In the top panel, we assume a zero percent discount rate. As shown in the “Total” row, we find that two-thirds (66.8%) of the valid observations elected to delay recognition to 2011 and 2012. Of these, we estimate that 44.2% (equal to $\frac{0.295}{0.668}$) of these individuals made the ex-post correct decision; this increases to 71.4% ($\frac{0.295+0.182}{0.668}$) when considering individuals for whom the same tax was owed under either scenario. Put differently, among those individuals for whom delaying was optimal, 75.3% ($\frac{0.295}{0.392}$) made the choice to delay. But among those for whom immediate recognition was optimal, only 41.2% ($\frac{0.133}{0.191}$) elected to recognize immediately.

The second panel assumes a 6 percent discount rate. As shown in the “Total” column, this generally has the effect of making delaying recognition more tax-favored – over 77% of taxpayers would pay less tax (in present value) by delaying recognition. Much of this increase (relative to 39.2% under the zero-discount-rate assumption) comes from the category of taxpayers who would pay the same amount of tax in nominal dollars under either decision – in this case, the time value of money under the 6 percent discount rate “breaks the tie.”⁶⁰ Mechanically, this increases the share of taxpayers who delay recognition that turn out to have made the correct decision ex-post – 84.4% ($\frac{0.558+0.006}{0.668}$), up from 71.4% under the assumption of a zero percent discount rate. This also mechanically decreases the “share correct” of taxpayers who recognize immediately – 36.1% ($\frac{0.099+0.021}{0.213}$) under a 6 percent discount rate rather than 70.8% ($\frac{0.133+0.102}{0.332}$) under a zero percent discount rate. Using the 6 percent discount rate, we find that the median extra liability among taxpayers who made an ex post suboptimal decision is \$320.⁶¹

These tables do not exploit variation in the magnitude of the benefit of delaying recognition. To illustrate the effect of this benefit in a more continuous manner, Figure 6 shows

⁵⁹The 2010 income recognition timing decision is much simpler than decisions associated with the lifetime allocation of retirement savings across pre- and post-tax accounts because the former only concerns expectations about relative marginal tax rates, while the latter involves all of those factors discussed in Section 1. Therefore, if a large fraction of individuals has difficulty with the simpler decision, it is reasonable to anticipate that at least as large a fraction will struggle to make an optimal decision in the more complicated case.

⁶⁰Some taxpayers – 2.7% – still pay the same amount of tax in present value under either decision. This could happen if the total amount of tax owed is zero under either decision (e.g., because the taxpayer does not have enough income, even with the Roth conversion, to exhaust personal exemptions and/or the standard deduction).

⁶¹Using a discount rate of zero, we find that the median extra tax liability is \$500.

a binned scatter plot of the probability of delaying recognition as a function of the absolute benefit of doing so. The four panels repeat the analysis at different assumed discount rates. In each of the panels, there is a broad upward relationship – suggesting that taxpayers have some foresight about the benefit of delaying recognition. However, this upward relationship is broken near zero. Intuitively, conversions which give rise to similar tax liability whether recognized immediately or delayed tend to be *small*, for two reasons. First, a smaller conversion amount is less likely to cause income to cross a bracket threshold – leading to constant nominal tax over time. Second, with a positive discount rate, the benefit of delaying is increasing in the conversion amount (holding marginal tax rates fixed). Thus, these figures also suggest that taxpayers who make small Roth conversions are less likely to delay recognition. One explanation for this finding is that such taxpayers are less sophisticated and would face a nontrivial hassle cost for delaying recognition – e.g., because it would require filing Form 8606 in 2011 and 2012 as well.⁶²

5 Characteristics of 2010 Roth Converters

In this section, we examine various characteristics of Roth converters over time. Of course, these characteristics changed mechanically due to the elimination of income limits on Roth conversions in 2010. Figure 7 plots the median income of Roth converters and the share of converters that is married over time. For example, among individuals that performed a Roth conversion in 2005, the median income was \$50,000 (2009 dollars) and approximately 64% of them were married. Both series increase substantially in 2010, and remain elevated thereafter: higher income tax filers, who are more likely to be married, were able to make a conversion due to the removal of the conversion income limits.

In Table 3, we compare covariates of Roth converters to the covariates of the ever-IRA sample as a whole, separately for the 2004-2009 period, 2010, and 2011-2015. All variables in this table (except age) are lagged by one year. This table shows that, indeed, converters tended to have lower income than non-converters prior to 2010, and higher income than non-converters in 2010 and 2011-2015.⁶³ Similarly, converters were less likely to be married and were slightly older in 2004-2009 than non-converters; this reverses in 2010 and 2011-2015.

There are some patterns that appear to hold in all periods. First, converters were more likely to be male. Second, converters have more savings, as measured by IRA balance,

⁶²In Appendix Figure A2, we plot a binned scatter plot of the probability of delaying recognition as a function of the size of the conversion amount. There is a clear upward relationship: the probability of delaying increases from just over 50% for the smallest conversions to nearly 80% for the largest. Furthermore, in unreported multivariate regression results, we find that the (log) conversion amount and the (inverse hyperbolic sine) benefit of delaying recognition are both important determinants of the probability of delaying recognition.

⁶³For this purpose, we define income to be adjusted gross income, less taxable IRA distributions (including Roth conversions), less capital gains included in AGI, plus the IRA deduction, plus pre-tax contributions to a workplace retirement plan, and plus the non-taxable portion of Social Security benefits.

investment income, and workplace qualified plan saving. Third, in all years, converters were more likely to exhibit Roth behavior prior to the conversion: they were 36 to 57 percentage points (depending on the period) more likely to hold a Roth IRA in the year prior to conversion, relative to the comparison group. They also had higher average contributions to workplace Roth qualified plans. Fourth, converters received a smaller share of their income in the form of wages, and were more likely to report some business income.⁶⁴

While these unconditional differences between converters and non-converters are interesting in their own right, they are somewhat difficult to interpret given the large differences in income and the somewhat smaller differences in age and the share married – each of which we expect to have large, direct effects on the other covariates in this table. To control for these differences, we compare Roth converters and non-converters after partialling out the effect of these three factors (income, age, and share married). More formally, let X denote the set of these three covariates. To explore differences along other dimensions y , such as saving behaviors and income composition, we need to partial out the effect of X . One approach would be to use a regression analysis: simply regress y on being a Roth converter, controlling for X . However, we expect the effects of age, income, and marital status on various outcomes to be highly interacted in nonlinear ways.

For this reason, we use a less parametric matching approach known as coarsened exact matching (CEM) (Iacus et al., 2009). Under this procedure, we “coarsen” (i.e., discretize) the space of X and place each individual i into a bin b_i . We assign a weight of one to all Roth converters. For everyone else in the sample (“non-converters”), we assign a bin-specific weight w_b to equalize the total weight of converters and the total weight of non-converters within each bin. This has the effect of creating a control group that matches the distribution across bins of Roth converters; as the discretization of X becomes increasingly fine, then the control group and the Roth converter group will converge to the same distribution of X . We drop any observations that are assigned to a bin with only Roth converters or only members of the control group – this is analogous to the “overlap” condition required for propensity score reweighting.

We implement this approach separately by year. For each year, we calculate 400 quantiles of lagged income within each year-of-birth/marital status cell.⁶⁵ Using these quantiles (and the individual’s year of birth and marital status), we assign each individual-year observation

⁶⁴For the sake of this table, we define business income as the presence of positive Schedule C income or positive Schedule E income.

⁶⁵We also lag marital status. In this procedure, we consider individuals who are married filing separately (MFS) to be non-married, since (like single filers) the reported income will correspond to the income of the individual, rather than the couple.

to bin b_{it} .⁶⁶ We drop nonfilers (whose marital status we cannot observe) and individuals with negative income. These restrictions together cause us to drop approximately 142,000 of the original 4.15 million Roth converter observations between 2004 and 2015.⁶⁷ Of the remaining 4.01 million converter observations, we are able to find a match (*i.e.*, the overlap condition is satisfied) for all but 26,000.

Figures 8 and 9 illustrate the effect of the CEM procedure in the case of 2010. The top panel of each figure plots percentiles of prior-year income (in 2009 dollars) or age for the converter group and the unweighted control group. The bottom panel of each figure uses the CEM weights. The top panel of Figure 8 shows the extent to which Roth converters have higher income than the control group, throughout the distribution; this distribution is essentially equalized in the bottom panel, using CEM weights. The top panel of Figure 9 shows the analogous plot for age. This shows that, while mean age is relatively similar between the two groups (as seen in Table 3), the distribution is quite different. In particular, converters have a smaller dispersion in age than non-converters. In other words, converters are more likely to be middle-aged than non-converters, while non-converters are more likely to be relatively young or relatively old. As in the case of income, the bottom panel shows that this distribution is essentially equalized when using the CEM weights. Thus, as intended, the CEM procedure creates a control group that matches very well along the target dimensions of prior income and year of birth.

We then use these CEM weights to repeat the comparisons from Table 3. These results are reported in Table 4. By construction, the share married, age, and income quantiles match between the two groups. Yet, this reweighting has only a modest effect on the differences in other covariates. In particular, the weighting has little effect on (lagged) Roth-specific outcomes (owning a Roth IRA and making Roth qualified plan contributions). This reveals an interesting, if unsurprising, fact: Roth converters appear to have revealed a prior affinity for Roth IRAs that cannot be explained by differences in age, income, or marital status. Additionally, it remains the case that Roth converters are more likely to be male, have higher savings (as proxied by IRA balance, investment income, and workplace qualified plan saving generally), and have non-wage sources of income.

⁶⁶Recall that our sample includes a balanced panel (aside from death and being younger than 18) of the full population of individuals that ever made a Roth IRA. However, when constructing the potential control sample for year t (*i.e.*, before re-weighting), we need this sample to be representative of the ever-IRA population that did not make a Roth conversion in year t . We do this by taking a 2.5% random sample (the same sampling rate as the ever-IRA sample generally) of the set of individuals who make a Roth conversion in some year but not year t .

⁶⁷Because this procedure uses lagged income and marital status, we cannot calculate CEM weights for observations in 2003.

6 Roth Conversion Motives

In this section, we consider several potential explanations for why individuals do Roth conversions. First, an individual might experience a transitive shock which reduces his or her marginal tax rate in a given year, making a Roth conversion relatively more attractive. Second, a Roth conversion can increase the amount of after-tax savings held within a tax-advantaged account. Third, an individual might wish to diversify against future rate uncertainty by holding some assets in Roth accounts and others in traditional account. Fourth, an individual might desire to take advantage of the fact that some Roth IRA assets can be withdrawn without tax or penalty even before age $59\frac{1}{2}$. Fifth, the salience of the 2010 policy event – along with the corresponding news coverage – may have caused some individuals to make a Roth conversion, either because of the rational analysis of the receipt of new information, or for behavioral reasons. We consider each motivation in turn. With the exception of gaining access to penalty-free distributions, we find suggestive evidence that each of these motivations play a role in encouraging Roth conversions.

6.1 Low Marginal Tax Rates

If an individual's marginal tax rate is transitively low in a given year, a Roth conversion is relatively more attractive, even if there is little uncertainty regarding future tax rates. In this subsection, we present descriptive evidence of the relationship between an individual's marginal tax rate and her probability of making a Roth conversion. The time series evidence is consistent with this explanation. In Figure 2, there are clear increases in Roth conversions in 2009 and 2012, in addition to the much larger spike in 2010. Due to the Great Recession, 2009 may have been a low-marginal-rate year for many individuals, as the business cycle caused income to fall, thus pushing people toward the lower end of the progressive rate schedule.⁶⁸ Additionally, 2012 was a low-tax year relative to future years due to the expiration of the Bush tax cuts for high-income individuals in 2013. We find that even lower income individuals responded in 2012 in anticipation of a possible rate increase in 2013, as the precise policy change was uncertain until January 1, 2013.

In Figure 10, we break down the increase in 2012 by income group (in 2009 dollars). We partition 2011 and 2012 converters into bins defined by income. The bins are set such that each bin contains (approximately) the same amount of Roth conversions in 2011, with the bin size equal to about \$400 million. The series with the solid line and solid circles plots the total amount of Roth conversions in 2012 in each of these bins.⁶⁹ For example individuals with 2011 income from \$110,000 to \$133,000 converted \$400 million in 2011; in 2012, this income group (defined based on 2012 income) converted \$600 million. The figure

⁶⁸In the next subsection, however, we will consider an alternative explanation for the increase in 2009.

⁶⁹The series for 2011 just reiterates the fact that the income bins were created to be the same size.

shows that Roth conversions increased in 2012 in all income bins with the ratio between 2012 conversions and 2011 conversions (nearly) monotonically increasing in income. This ratio is approximately three for the highest income bin, which includes individuals with income in excess of about \$2.4 million: that is, among individuals with income greater than \$2.4 million, aggregate conversions were 3 times larger in 2012 relative to 2011.

It is tempting to try to exploit variation in tax changes between 2012 and 2013 to recover a causal estimate of future tax changes on present Roth conversion behavior in a difference-in-differences framework. However, such an approach would be complicated by the fact that the parameters of the rate increase between 2012 and 2013 were not known until early January 2013. In essence, the “control” group which did not experience a rate change might have acted in 2012 as if it were treated, attenuating a difference-in-differences estimate. Instead, to explore our hypothesis further, we use a distributed lag framework, which shows the path of Roth conversions around marginal tax rate changes experienced by individuals, whether caused by policy or income changes, where we define the marginal tax rate to be the rate that applies to the first dollar of IRA distributions.⁷⁰ Specifically, let y_{it} denote a dummy for making a Roth conversion in year t , and let mtr_{it} denote the marginal tax rate (in percentage points) facing the first dollar of IRA distributions. We are interested in the response of y_{it} to changes in the net-of-tax rate, $1 - mtr_{it}$. We run the following regression, using a specification in first-differences.⁷¹

$$\Delta y_{it} = \delta_t + \sum_{k=-2}^2 \tilde{\beta}_k \Delta(1 - mtr_{i,t-k}) + \epsilon_{it} \quad (2)$$

Suppose there is a transitory, one percentage point decrease in the marginal tax rate at time τ (i.e., an increase in the net-of-tax rate), reverting back to the previous value at time $\tau + 1$. The $\tilde{\beta}_k$ coefficients will trace out the effect on the probability of making a Roth conversion between $\tau - 2$ and $\tau + 2$. For the sake of interpretation, we normalize these effects relative to $\tau - 1$; i.e., we report $\beta_k = \tilde{\beta}_k - \tilde{\beta}_{-1}$. These coefficients are plotted in the left panel of Figure 11; the series with the solid line and solid circles omits time fixed effects, while the series with the dotted line and hollow circles includes them.⁷² These plots show that a transitory one percentage point increase in the net-of-tax rate leads to an increase in Roth conversions by about 0.01 percentage points, from a baseline of 0.40 percent. Put differently,

⁷⁰We use the same tax calculator described in Section 3.2. If taxpayers fill out Form 1040 correctly, income attributable to a Roth conversion will be recorded on line 15b (of the pre-2018 Form 1040) with all other IRA distributions. Thus, the marginal tax rate will not be *mechanically* affected by Roth conversions.

⁷¹If estimated in levels, the earliest lead ($k = -2$) and latest lag ($k = 2$) would include effects from periods earlier than $k = -2$ and later than $k = 2$. In first differences, the earliest lead and latest lag includes effects only from those specific years.

⁷²These coefficients, and those discussed in the next paragraph, are also reported in Appendix Table A1. That table also reports the associated standard errors. All variables are significant, except the second lead for the transitory tax change and the first lag for the transitory tax change in the specification with time fixed effects.

a transitory ten percentage point reduction in the marginal tax rate would increase Roth conversions by about 25% of the baseline amount of conversions. This is consistent with taxpayers responding to marginal rate changes by adjusting their Roth conversion behavior. The conversion probability returns close to baseline in the year after the change, suggesting that the increase in conversions is not a timing shift.

We can also think about these results in the context of a one percentage point persistent increase in the net-of-tax rate at time τ . The estimate for the cumulative effect of this change at time $\tau - k$ (relative to $\tau - 3$), denoted by $\tilde{\rho}_k$, corresponds to the sum of the $\beta_{k'}$ coefficients from $k' = -2$ to $k' = k$. We again normalize these cumulative effects relative to $k = -1$; i.e., $\rho_k = \tilde{\rho}_k - \tilde{\rho}_{-1}$. We plot these results in the right panel of Figure 11; as above, the series with the solid line and solid circles omits time fixed effects, while the series with the dotted line and hollow circles includes them. Under this specification, a persistent increase in the net-of-tax rate leads to a persistent increase in Roth conversions. Thus, our evidence is consistent with taxpayers responding to a decrease in marginal tax rates by increasing Roth conversions, though it is not causal.

These results are not entirely consistent with the spike in conversions seen in 2012. If the predominant tax effect of the expiration of EGTRRA and JGTRRA in 2013 is a persistent increase in marginal tax rates (or decrease in net-of-tax rates) in 2013 relative to 2011 and 2012, then these coefficient results would predict a roughly flat path of Roth conversions in 2011 and 2012, followed by a reduction in 2013 which is maintained in 2014 and 2015. In other words, we should expect to see Roth conversions follow the opposite of the right panel of Figure 11, with $t = 0$ corresponding to 2013. The fact that we actually see a spike in 2012 relative to 2011 suggests that high-income individuals – who are largely responsible for the 2012 spike – behave somewhat differently than Roth converters as a whole.

6.2 Increase Tax-Advantaged Saving

A Roth conversion can potentially increase the amount of savings held within a tax-advantaged account. This is most obvious in the case of “backdoor” conversions. A backdoor conversion is economically equivalent to a direct contribution to a Roth IRA, which increases saving in a tax-advantaged account mechanically, unless it is offset by a reduction in other tax-advantaged saving. Thus, backdoor conversions might be attractive to those individuals who cannot increase their retirement saving at work, through deductible IRA contributions, or through direct Roth IRA contributions.

More subtly, even a taxable Roth conversion (i.e., a conversion where basis is zero or small) can increase tax-advantaged saving: the tax paid on the Roth conversion is effectively money shifted from a taxable to tax-deferred account. Consider the following simple example. Suppose an individual holds \$100,000 in a traditional IRA and \$25,000 in a taxable account. She faces a constant, certain tax rate of 25 percent in all years, which applies equally to IRA

distributions and investment income in the taxable accounts.⁷³ Assume that she will receive a certain, annual return on her assets r between now and retirement T years in the future; r is assumed identical between the taxable account and the IRA. If she engages in a Roth conversion, she pays \$25,000 in (extra) tax this year, exhausting her taxable account. When she retires, she will be able to consume $\$100,000(1+r)^T$ from her Roth IRA. So, her total consumption will be $\$100,000(1+r)^T$.

If she does not do a Roth conversion, her taxable account will be worth $\$25,000(1+r(1-0.25))^T$ and her IRA will be worth, after tax, $\$100,000(1+r)^T(1-0.25)$. Her total consumption will be $\$25,000(1+r(1-0.25))^T + \$100,000(1+r)^T(1-0.25)$. Comparing consumption in retirement under a Roth conversion to consumption without a Roth conversion, one finds that the former is higher than the latter, as follows:

$$\begin{aligned}
C^{conv} - C^{noconv} &= \$100,000(1+r)^T - & (3) \\
& (\$25,000(1+r(1-0.25))^T + \$100,000(1+r)^T(1-0.25)) \\
& = \$25,000(1+r)^T - \$25,000(1+r(1-0.25))^T \\
& > 0
\end{aligned}$$

That is, under a Roth conversion, the \$25,000 tax paid on the initial conversion earns a higher rate of return (r) than it would have had it remained in the taxable account ($r(1-0.25)$). This suggests that taxable Roth conversions might be attractive for those individuals who are looking to save more in tax-advantaged accounts, but who cannot through other means.

In Table 5, we explore the statutory constraints faced by converters (and subsets thereof) and a control group.⁷⁴ In the first two columns, we use the coarsened-exact-matching (CEM) weights, as described in Section 5, to compare converters to a control group that matches the distribution of (lagged) income, marital status, and age. In the first row, we report the share of individuals who appear to not have access to a workplace retirement plan.⁷⁵ Approximately 29% of individuals in both the converter and matched control groups are ineligible to participate in an employer provided plan. In the second row, we calculate the (unconditional) share of individuals who are contributing the maximum amount allowed to a workplace plan.⁷⁶ Among individuals in the converter group, 22.3% are contributing the maximum, compared to individuals in the matched control group, among whom 12.9% are

⁷³For simplicity, assume that this is taxed on accrual, e.g., like interest income.

⁷⁴Because we are interested in saving behavior prior to retirement, we restrict the sample to those under age 55.

⁷⁵We classify employer j as having a workplace qualified (DC) plan at time t if we identify at least one worker (in the underlying universe of Forms W2) working for firm j at time t with a DC contribution indicated on their W2. We assume that all individuals who work for firm j at time t are eligible to participate in that plan.

⁷⁶We classify a contribution to be “at the maximum” if it is greater than 90 percent of the statutory maximum.

contributing the maximum.⁷⁷ The third row reports the sum of the first two rows, which represents the share of individuals who are unable to increase their retirement savings at work – either because the workplace does not offer a DC retirement plan, or because the individual is already contributing at the maximum. We find that 51.3% of converters are constrained, compared to 41.8% of the matched control group.

The bottom panel repeats the analysis, broadening the universe of retirement saving to include contributions to Roth IRAs and deductible contributions to traditional IRAs. The first row of this panel reports the share of individuals who are ineligible to make contributions to workplace DC retirement plans (because the firm does not offer one) and are ineligible to make Roth IRA contributions (due to income limits and/or lack of taxable compensation) and are ineligible to make deductible IRA contributions. These include individuals who, for instance, have high income and participate in their firm’s defined benefit retirement plan (making them ineligible for deductible IRA and Roth IRA contributions), but their firm does not offer a defined contribution plan. It also includes individuals with high income but no taxable compensation. We find that 14% of converters and 12% of the matched control group fit this description. In the second row of this panel, we report the share of individuals contributing at the maximum. For those eligible for either IRAs or workplace DC plans (but not both), we classify individuals as contributing at the maximum if they are contributing at the maximum for the account for which they are eligible. If they are eligible for both IRAs and workplace DC plans, we require individuals to be contributing at the maximum for both types of accounts. We find that the share of individuals contributing at the maximum across the workplace retirement and IRA plans for which they qualify is quite similar to that in the top panel: 23.5% for converters and 13% for the matched control group. The final row presents the sum of those ineligible and those contributing at the maximum, which again reflects the share that is constrained. We calculate this to be 37.7% among converters and 24.8% among individuals in the matched control group.

Roth converters are more likely to be constrained in making retirement savings contributions through other means, relative to the matched control group. This is consistent with such constraints motivating Roth conversions. However, this pattern could also be explained by Roth converters having a higher desire to save in any form, conditional on income, as observed in Table 4.

The third and fourth columns of Table 5 provide stronger evidence that retirement saving constraints are important for backdoor converters in particular. In these two columns, we repeat the same calculations for backdoor and taxable converters, restricting the sample to 2010 onwards.⁷⁸ We find that 48% of backdoor converters contribute the maximum amount

⁷⁷Note that this implies that the share of converters contributing at the maximum, conditional on being eligible, is $\frac{0.223}{1-0.290} \approx 0.314$.

⁷⁸We define a backdoor converter to be a converter with basis of at least 80% of the conversion amount; converters with basis below this threshold are classified as taxable converters.

to their workplace DC plan; an additional 22% have no access to such a plan, for a total share constrained of 71%. This figure drops modestly to 61.5% when including IRAs. In comparison, among individuals in the taxable converter group, 49% are constrained in contributions to workplace retirement plans and 35% are constrained when account for workplace plans and the IRAs for which they are eligible to contribute. That backdoor converters are 78% more likely to be constrained relative to taxable converters, using eligibility for both employer-provided and IRA plans, might be explained by the fact that the link between maximizing total tax-advantaged saving and Roth conversions is far more direct for backdoor converters than for taxable converters.

6.3 Diversification Against Future Tax Rate Uncertainty

Future tax rates are uncertain, both because an individual's own income is uncertain and because future tax policy is uncertain. For this reason, some individuals may find it optimal to hold some mix of assets with traditional tax treatment (i.e., upfront deduction with taxable distributions) and assets with Roth tax treatment (no upfront deduction, but tax-free distributions). As a first piece of evidence that diversification may motivate Roth conversions, we plot the share of taxable Roth conversions that involve a complete conversion of all IRA assets (excluding SEP and SIMPLE IRAs).⁷⁹ This figure shows that between 30% and 55% of Roth converters, depending on the year, did not convert their entire traditional IRA account, suggesting that these taxpayers used Roth conversions as a tax diversification strategy.

There is another explanation for why some individuals might not convert the entirety of their IRA assets in a single year: repeated conversions over many years can spread the income recognition from those conversions, which may be optimal given the progressive rate structure of the individual income tax. In Figure 13, we investigate this explanation. Among 2010 converters who did not convert the entirety of their IRA assets in 2010, we plot the share of individuals who have zero traditional IRA assets at the end of the year. This share is zero in 2010 by construction. By 2011, about 7.7% of less-than-full converters had reduced their traditional IRA assets to zero, either through a distribution or a subsequent Roth conversion. The share increases each year but only to about 15% by the end of 2015. This suggests that relatively few less-than-full converters engaged in a “slow motion” full conversion over many years, which is consistent with the use of Roth conversions as a tax diversification strategy.⁸⁰

⁷⁹For this exercise, we disregard conversions with basis share of at least 80 percent (i.e., backdoor conversions), since those do not represent a conversion from traditional tax treatment to Roth tax treatment.

⁸⁰It is also possible that a large number of individuals plan to slowly convert all traditional assets over more than 6 years, in which case we would underestimate the total fraction of full converters over these years and overstate the case for the diversification motive. However, the concavity of the graph indicates that conditional on not making a full conversion in the previous year, an individual is less likely to do so with each year that passes.

Most individuals with both pre- and after-tax accounts hold the majority of their IRA assets in traditional accounts. In Figure 14, we plot the empirical density of the Roth share – that is, the account balance of the Roth IRA divided by the total of Roth and traditional IRA balances – in 2011 and 2015, among the set of individuals with both types of accounts.⁸¹ We find that the density is nearly monotonically decreasing, meaning that most of these dual-account holders are holding most of their assets in traditional IRAs. Furthermore, this empirical density is quite stable since 2011, as the densities in 2011 and 2015 are nearly indistinguishable. In particular, the median Roth share is 20% in each year between 2011 and 2015.

There is a slight bump evident at exactly 50%. This may be a data artifact. In particular, suppose an individual receives a 5498 for a Roth IRA account with FMV V which erroneously indicates that the IRA is a traditional IRA. Suppose that individual later receives a corrected 5498, listing the same value V , in which the IRA is correctly labeled as a Roth IRA. Our data-cleaning algorithms would not know that the original 5498 is obsolete. Both Forms 5498 would show up in our data and we would erroneously calculate a Roth share of 0.5. On the other hand, it is also possible individuals use a 50% rule-of-thumb when deciding how to allocate their assets across accounts, as diversification biases tend to bias savings choices toward an equal allocation (Fox et al., 2005; Beshears et al., 2017).

6.4 Penalty-Free Distributions

In general, distributions from a traditional IRA prior to age $59\frac{1}{2}$ incur a 10% penalty (in addition to being subject to income tax, which is true regardless of age). By contrast, individuals may generally take a distribution from a Roth IRA up to his or her basis in the account – which is equal to contributions plus conversions plus rollovers, less any previous distributions – without penalty. This creates a potential motive for a Roth conversion: individuals can gain penalty-free access to the amounts in the account prior to retirement. However, the Internal Revenue Code requires a five-year waiting period after a conversion before the amount of the conversion can be withdrawn without penalty. For instance, if an individual made a conversion in 2008, any distributions prior to 2013 would be subject to the 10% penalty; beginning in 2013, distributions up to the amount converted may be taken without penalty.⁸²

If this motive plays a large role in motivating Roth conversions, we should see an increase in distributions from Roth IRAs five years after a conversion. To test this, we estimate an event-study regression among the set of individuals who make a conversion between 2005

⁸¹We disregard SEP and SIMPLE IRAs for this purpose.

⁸²In this example, we assume that no other contributions are made to the Roth IRA, either before or after the conversion. More generally, the Internal Revenue Code specifies an ordering rule for determining which assets are being distributed in which order for tax purposes. See IRS Publication 590-B for additional details.

and 2010. We construct a balanced eight-year panel (from two years prior to the conversion to five years after) for each person/conversion-year observation. The regression equation is as follows, where k indexes event time (i.e., years since conversion), τ indexes the year of the conversion, and $y_{it\tau}$ is a dummy for person i taking a “regular” Roth IRA distribution (with 1099-R distribution codes J, Q, or T) in calendar year t .

$$y_{it\tau} = \alpha_t + \sum_{k=-2, k \neq -1}^5 \beta_k \times \mathbb{1}(t - \tau = k) + \epsilon_{it\tau} \quad (4)$$

The β_k coefficients measure the rate at which individuals are taking Roth IRA distributions at event-time k , relative to $k = -1$, controlling for calendar-year fixed effects. We plot these results in Figure 15; for the sake of interpretation, we vertically shift the graph by the baseline value at $k = -1$. I.e., about 1.8 percent of individuals take a Roth IRA distribution at $k = -1$, corresponding to the value reported on the graph. All other points correspond to this 1.8 percent baseline level plus the corresponding β_k coefficient.⁸³ This figure shows that, unsurprisingly, distributions increase in the year of a conversion, as many converters do not have a Roth IRA at all prior to the conversion. For $k > 0$, there is a relatively smooth, increasing relationship between event-time and the probability of taking a Roth IRA distribution. There may be an inflection point at $k = 5$, the first year when distributions can be made out of the conversion amount penalty-free, though it is very slight and is difficult to distinguish from the pattern over the prior several years. Thus, we do not find substantial evidence that gaining access to IRA assets is a major motive for Roth conversions.⁸⁴

Our finding differs from recent research which shows a sizable response to the 10% penalty: Goda et al. (2018) find that among individuals around the $59\frac{1}{2}$ age threshold, there is an 80% increase in annual withdrawals upon reaching the age cutoff. Additional research shows that leakage from retirement accounts in general is substantial: approximately \$0.40 of every \$1 contributed into an account prior to age 55 is withdrawn from the account prior to retirement (Bryant et al., 2010; Argento et al., 2015). Therefore, we were surprised that there is little evidence of a response in the case of conversions. It is possible that individuals who are aged near the penalty age threshold respond differently to the penalty compared to individuals in our much wider age group, for example because of salience. The $59\frac{1}{2}$ age threshold is likely more well-known relative to the 5-year rule studied here. Additionally, it is possible that Roth converters are less sensitive to the penalty because they

⁸³The figure also reports 95 percent confidence intervals, which are extremely tight.

⁸⁴We also implemented this procedure using a panel from two years prior to six years after a conversion. Because there is empirical evidence of taxpayer confusion regarding IRA rules in general, it seemed possible that individuals would wait until year 6, instead of year 5, to take a penalty-free distribution. However, the results from that are consistent with those presented here: that is, we do not find evidence that Roth distributions increase markedly in the sixth year following a conversion. We choose to report the 5-year results rather than the 6-year results because the 6-year specification requires dropping 2010 converters.

have a stronger savings preference relative to non-converter IRA-holders, as shown in this paper. This hypothesis is also consistent with Mortenson et al. (2017), who find that Roth IRA withdrawals are much smaller compared to withdrawals from traditional IRAs among retirement-age individuals.

6.5 Information & Confusion

Lastly, it is possible that some individuals responded to new information about Roth conversions in the years leading up to 2010. This response may or may not have been optimal for the individuals involved. We discussed the first piece of evidence in favor of this hypothesis in Section 4.3. In that section, we found that approximately 30% of 2010 Roth converters made an ex-post suboptimal decision regarding whether to recognize the income immediately, or in 2011 and 2012. Given that the prediction of future tax rates is close to the heart of the Roth conversion decision, the fact that 30% of converters could not correctly predict their rate two years into the future suggests that at least as many individuals will incorrectly predict their marginal rate further into the future.

We present another piece of evidence in favor of some informational effect in Figure 16. The figure plots the number of Roth converters with MAGI greater than \$110,000 (solid line) and less than \$90,000 (dotted line), by year. Prior to 2010, essentially zero individuals with MAGI greater than \$110,000 engaged in a Roth conversion, as they were above the \$100,000 limit; this series spikes in 2010 when these limits are repealed, and drops down to a baseline level in the following years. More interestingly, we also see a nearly 100% increase in 2010 (relative to 2009) in the number of converters with MAGI less than \$90,000. These individuals should not have been affected by the elimination of the income thresholds, as they could make conversions prior to 2010, yet there is a substantial response among this group.

There is another candidate explanation for the spike in 2010 conversions among those with MAGI under \$90,000: the ability to delay the recognition of their Roth conversion income to 2011 and 2012. This option was not available to converters in other years, creating an incentive to perform the conversion in 2010 rather than an earlier year.⁸⁵ Thus, we cannot reliably conclude that the spike in conversions in 2010 among those with income less than \$90,000 is evidence of confusion or information. On the other hand, the increase in conversions among lower-income individuals in Figure 16 began in 2009 – when there was no option to delay recognition to later years. This is consistent with a response to news stories in advance of the elimination of the income limit.⁸⁶

We present another puzzle which is hard to explain with neoclassical explanations. Gen-

⁸⁵We find that about 55% of converters with income under \$90,000 who fit the sample restrictions in Section 4.3 choose to delay recognition to 2011/2012.

⁸⁶The 2009 response could also be the result of changes in marginal tax rates, as discussed previously.

erally, Roth converters reveal a preference to hold their assets in Roth-type accounts. Yet, a large number of these individuals did not make Roth contributions when they had an option to do so. In the first row of Table 6, we report the share of converters and non-converters who make a contribution to a Roth qualified plan (e.g., 401(k)) at work.⁸⁷ We find that 7.1% of Roth converters make a Roth contribution at work, compared to 1.3% for the rest of our sample. In the second row, we apply the CEM weights, as discussed in Section 5, to control for differences in income which may be correlated with differences in contributions to Roth employer-provided retirement plans. This changes the value among Roth converters only slightly, and increases the share of the control sample making Roth contributions from 1.3% to 3.1%.⁸⁸ Thus, controlling flexibly for differences in income reduces but does not eliminate the gap in Roth contributions. Nevertheless, only 7% to 8% of Roth converters made Roth contributions to workplace retirement plans, despite their apparent preference for Roth treatment rather than traditional treatment. The most likely explanation for this low share is that not all employers allow for Roth contributions to their plan. The next row in Table 6 restricts the sample to those who appear to work for a firm that allows Roth contributions.⁸⁹ This increases the share of Roth converters who made a Roth contribution at work to 0.23. However, this share is far less than 1, which suggests that over three quarters of Roth converters do not take advantage of Roth options at work.⁹⁰ This is further evidence that there may be some behavioral or informational component to the Roth conversion decision.⁹¹

Finally, we show that a requirement that is unrelated to Roth conversions affects them, suggesting a role for hassle costs or confusion. Individuals must take RMDs from their traditional IRA beginning in the year in which the individual turns $70\frac{1}{2}$: the amount of the RMD is equal to a percentage of the balance at the end of the prior year and is increasing in age. It is theoretically ambiguous whether RMD eligibility should affect conversions. On the one hand, while there is no prohibition on making a Roth conversion after age $70\frac{1}{2}$, taxpayers

⁸⁷Throughout this table, we include spouses' contributions and eligibility.

⁸⁸Recall that any changes in the figure for converters reflects minor sample restrictions, such as the restriction to have nonnegative income, as well as the restriction to the set of converters that satisfy the overlap condition.

⁸⁹We classify employer j as having a Roth option at time t if we identify at least one worker (in the underlying universe of Forms W2) working for firm j at time t with a Roth contribution indicated on their W2.

⁹⁰We do not report the analogous figure using the matched control sample. Conditioning on availability of Roth workplace accounts would cause the matched control sample to no longer match the income, age, and marital status distribution of the converter sample.

⁹¹It is possible that Roth converters have a particular preference for after-tax savings in IRAs, as opposed to employer-provided plans. Defined contribution (DC) plans with a Roth option are different from Roth IRAs in several ways: DC plans do not have income limits associated with contributions, have much higher contribution limits, and are subject to RMDs. They also differ in other ways that tend to depend on the specific plan: for example, the rules associated with borrowing from the accounts differ. However, we are surprised at the magnitude of the share of Roth converters with access to Roth DC plans that do not take advantage of them and find it unlikely that a difference in preferences for the two types of accounts drives the entire result.

may find RMD rules – and their interaction with Roth conversion rules – to be confusing.⁹² In order to withdraw a required minimum and make a Roth conversion in the same year, the total distribution must exceed the conversion amount by at least the RMD – e.g., one could take the RMD out of the IRA and convert the rest. This extra transaction may cause sufficient marginal hassle to make the Roth conversion not worthwhile. On the other hand, the RMD could actually encourage a Roth conversion, for two reasons. First, the RMD itself can help pay for the tax liability created by the Roth conversion. Second, Roth IRAs do not require RMDs; thus, a Roth conversion may be attractive to an individual who wishes to avoid RMDs in the future.⁹³

To determine the effect of RMDs on Roth conversions, we exploit the discontinuity in exact date of birth. Individuals that turn $70\frac{1}{2}$ in December of year t must make an RMD for that year.⁹⁴ Individuals that turn $70\frac{1}{2}$ in January of the year $t+1$ are not required to take an RMD for year t . Figure 17 plots the converter share in 7-day date-of-birth bins (conditional on having a positive IRA balance in the prior year). Those on the left side of zero (indicated by the dotted line) do not attain $70\frac{1}{2}$ by the end of the year and are not required to take an RMD for year t ; those to right are slightly older, and thus are required to take an RMD for year t . The figure shows that the discontinuity in the outcome is sharp and large: the share making a conversion drops from approximately 0.75 percent to 0.58 percent – a drop of over 20%. Appendix Table A2 shows the associated regression discontinuity estimates with a variety of bandwidths. The coefficients are relatively stable between 0.0017 and 0.0019 and are highly significant, with the smallest t -statistic in excess of 3.5.⁹⁵ In sum, we find robust evidence that RMD eligibility reduces Roth conversions, which can most easily be explained by assuming that many older taxpayers find the rules regarding IRA contributions, Roth

⁹²There is no age limit associated with Roth contributions. Individuals can also make rollover contributions to traditional and Roth IRAs regardless of age. The IRS Publications 590-A and 590-B contain more details about IRA rules.

⁹³Because individuals cannot make traditional IRA contributions after age $70\frac{1}{2}$, only individuals who have an existing traditional IRA, or create a new one because of a rollover from a defined contribution account, are able to make a conversion after the age. Therefore, the sample of individuals included in this analysis is limited to those who have a traditional account in the prior year. Consistent with the tax code, this sample restriction eliminates individuals who would like to make a backdoor Roth contribution but cannot because they cannot make a non-deductible contribution to a new traditional IRA.

⁹⁴Individuals who turn $70\frac{1}{2}$ in a given year are given a grace period until April 1 the following year to make this distribution, potentially reducing the implicit “first stage” of this regression.

⁹⁵The key regression discontinuity assumption for this parameter to represent a causal effect is that there is no sorting around the cutoff point. In the present context, this requires that mothers are not choosing precisely to give birth just before or just after the end of June, which we view as a reasonable assumption. In Appendix Figure A3, we plot the number of individuals in 7-day bins and, indeed, find it to be smooth around this threshold. Additionally, it is possible that this discontinuity affects other outcomes (e.g., the age of entry into school) which could have downstream effects on Roth conversions. To allay this concern, we repeat the same analysis as Figure 17 using the placebo threshold of age $68\frac{1}{2}$; any spurious effects of exact date-of-birth would presumably have a similar effect at age 68.5 as age 70.5 These results are shown in Appendix Figure A4; reassuringly, we see no discontinuity in Roth conversion probabilities at this placebo threshold.

conversions, and RMDs – and the interactions thereof – to be confusing.

7 Conclusion

We study Roth conversion behavior from 2003 through 2015. The income limits associated with Roth conversions were removed in 2010, though there continue to be income limits associated with Roth contributions. The policy change allowed high-income individuals to convert traditional IRA assets which had accumulated over prior years. In the years since, it has enabled individuals to perform annual “backdoor Roth” contributions, which approximate direct contributions. The extent to which previously Roth contribution-ineligible individuals responded to, and continue to respond to, the 2010 policy change has remained an open question.

We use a unique source of data to study responses to the removal of the conversion income limits in 2010 and document the prevalence of backdoor Roth contributions in the years thereafter. Using administrative tax data on the universe of individuals who made conversions from 2003 through 2015, we show that there was a large response to the policy change in 2010: aggregate conversions increased by over 700% compared to the prior year. Furthermore, an increasing share of converters are backdoor converters: while approximately 1.5% of converters were backdoor prior to the policy change, by 2015 more than 37% of Roth converters were backdoor converters.

Our data, the 2010 policy change, and variation in Federal tax policy allow us to explore several reasons why individuals prefer to make after-tax instead of pre-tax retirement savings contributions. Specifically, we find supporting evidence for several motives, some of which have been documented in other studies: individuals are more likely to convert in a transitory low tax bracket year, use Roth accounts as a form of diversification against future tax rate uncertainty, and, among individuals who contribute the maximum to their other retirement savings accounts, use conversions to increase after-tax savings. We do not find evidence indicative of individuals being motivated to make a conversion because they plan to make a penalty-free withdrawal (of the conversion amount) from the Roth after the 5-year waiting period, which individuals are not permitted to make from traditional accounts.⁹⁶

For Federal tax purposes, the income associated with a conversion is included in income in the year in which the conversion occurs. However, individuals who made conversions in 2010 could elect to recognize the resulting income immediately in 2010, or in equal amounts in 2011 and 2012. We find that approximately 67% of individuals chose to delay recognition. Using a simplified tax calculator, we estimate that approximately 31% of individuals made the ex post incorrect decision: they could have reduced the tax liability associated with the conversion by making the opposite choice. Of these 31%, we estimate that the median

⁹⁶Mortenson et al. (2017) provide evidence of an additional motive: individuals may be motivated to save in Roth IRAs to avoid the RMD rules associated with traditional accounts.

of the extra tax liability was \$320. That a non-trivial fraction of individuals experienced difficulty making an ex post optimal decision over such a short time is indicative of the difficulty individuals face in making optimal retirement savings decisions over the course of time between when they make contributions and draw down assets, which could span 30 or 40 years. We also present evidence that suggests information, and possibly confusion about tax policy, play a role in retirement savings decisions.

Despite the sizable responses, Roth conversions are the exception rather than the rule among IRA holders: only 3.3% of IRA-holders over the 2003-2015 period made a Roth conversion. This number suggests a relatively low take-up rate of conversions among newly eligible individuals following the 2010 policy change. Conditional on having a traditional or Roth IRA in the prior year, only 0.3% of individuals made a conversion prior to 2010, 1.7% made a conversion in 2010, and between 0.7-0.9% made a conversion annually from 2011 to 2015.⁹⁷ Among the group of individuals most likely to make a conversion – those whose income was too high for a direct Roth contribution – the share who made a backdoor Roth contribution after 2010 was comparatively much higher: 0.7% made a backdoor Roth contribution in 2010 and by 2015, 1.6% of individuals made a backdoor Roth contribution. In aggregate, the use of conversions is quite low and we document the characteristics of Roth converters, a group which represents a small fraction of taxpayers.

There are several potential reasons why the use of conversions is low. First, for taxpayers with sizable assets in traditional IRAs, the taxable income associated with a conversion could be quite large and cause a taxpayer to be in a higher tax bracket than they would have been without the conversion. Accounting for basis across all traditional IRAs likely serves as a deterrent to executing a conversion for some taxpayers and we view this as the most likely explanation for why take-up of conversions has been low following the 2010 policy change.⁹⁸ However, basis represented very little of the 2010 spike in total conversions when measured in dollars: among those who made conversions, basis was low. Therefore, low basis was not a deterrent to converting for those who converted. Alternatively, there have been concerns about the legality of the backdoor Roth IRA. Specifically, the IRS Step-Transaction doctrine states that you cannot take multiple legal steps that, when taken together, constitute an illegal action. We do not have any evidence as to how many individuals would like to make Roth contributions, are unable to because of the direct contribution income limits, and are unwilling to make a backdoor Roth contribution because of legality concerns. Lastly, it could be that demand for after-tax savings is low among individuals with income that is too high to make a direct Roth contribution. This seems unlikely, given that after-tax savings relative

⁹⁷Because these figures pertain to individuals who had any type of IRA in the prior year, the population averages are much smaller.

⁹⁸It is possible that in anticipation of the 2010 policy change, financially sophisticated individuals did not rollover pre-tax assets saved in defined contribution plans into traditional IRAs so that they could maximize their basis in traditional IRAs and take advantage of the policy change at minimal cost. We have not explored this empirically.

to pre-tax tends to increase with income and, by our calculations, aggregate deductible contributions to traditional IRAs were \$10.5 billion and aggregate direct contributions to Roth IRAs were \$21.5 billion in 2015.⁹⁹ Roth contributions are much larger in aggregate compared to deductible traditional IRA contributions, especially among younger savers who have more of an incentive to save in Roth IRAs. In sum, our findings show there is substantial heterogeneity across taxpayers in demand and motivation for pre- and post-tax retirement savings accounts.

References

- Argento, Robert, Victoria L. Bryant, and John Sabelhaus (2015), “Early Withdrawals from Retirement Accounts During the Great Recession.” *Contemporary Economic Policy*, 33, 1–16.
- Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian (2014), “Who Uses The Roth 401(k), and How Do They Use It?” In *Discoveries in the Economics of Aging*, 411–440, University of Chicago Press.
- Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian (2017), “Does Front-Loading Taxation Increase Savings? Evidence From Roth 401(k) Introductions.” *Journal of Public Economics*, 151, 84–95.
- Bryant, Victoria L. and Jon Goyer (2013), “Accumulation and Distribution of Individual Retirement Arrangements, 2010.” *Statistics of Income Bulletin – Fall 2013*.
- Bryant, Victoria L., Sarah Holden, and John Sabelhaus (2010), “Qualified Retirement Plans: Analysis of Distribution and Rollover Activity.” In *Proceedings. Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association*, volume 103, 81–90, National Tax Association.
- Burman, Leonard E. (2006), “Roth Conversions as Revenue Raisers: Smoke and Mirrors.” *Tax Notes – May 22, 2006*.

⁹⁹The aggregate basis of conversions was approximately \$2 billion in 2015. Total assets held in traditional IRAs were \$6.4 trillion and total Roth assets were \$620 billion.

- Cilke, James (2014), “The Case of The Missing Strangers: What We Know and Don’t Know About Non-Filers.” Working Paper.
- Fox, Craig R., Rebecca K. Ratner, and Daniel S. Lieb (2005), “How Subjective Grouping of Options Influences Choice and Allocation: Diversification Bias and the Phenomenon of Partition Dependence.” *Journal of Experimental Psychology: General*, 134, 538–551.
- Goda, Gopi Shah, Damon Jones, and Shanthi Ramnath (2018), “How Do Distributions from Retirement Accounts Respond to Early Withdrawal Penalties? Evidence from Administrative Tax Returns.” Working Paper.
- Hoffman, Michael J.R. (2007), “Income Limits on Roth Conversions to Be Lifted in 2010.” *Journal of Retirement Planning*, 85, 47.
- Iacus, Stefano, Gary King, and Giuseppe Porro (2009), “CEM: Coarsened Exact Matching Software.”
- Kotlikoff, Laurence J., Ben Marx, and David Rapson (2008), “To Roth or Not? – That Is The Question.” Working Paper 13763, National Bureau of Economic Research.
- Mortenson, Jake, Heidi Schramm, and Andrew Whitten (2017), “The Effects of Required Minimum Distribution Rules on Withdrawals from Traditional IRAs.” Working Paper.
- Plan Sponsor Council of America (2012), “PSCA Annual Survey Shows Company Contributions Are Bouncing Back.” Press Release.
- Saez, Emmanuel (2017), “Taxing the Rich More: Preliminary Evidence from the 2013 Tax Increase.” *Tax Policy and the Economy*, 31, 71–120.

Tables

Table 1: Summary statistics

	(1)	(2)	(3)
	Population	IRA sample	Roth converter sample
Share male	0.488	0.495	0.525
Share married	0.567	0.672	0.704
Age	47.1	52.6	48.4
Share age 65+	0.190	0.245	0.159
Income-eligible for IRA deduction	0.892	0.804	0.676
Share with an IRA	0.189	0.533	0.533
Share with a Roth IRA	0.069	0.194	0.656
Eligible for Roth conversion (pre-2010)	0.874	0.756	0.618
Share ever making a Roth conversion	0.012	0.033	1.000
Share filers	0.809	0.911	0.975
AGI: 25th percentile	4,900	23,200	42,700
AGI: 50th percentile	28,500	58,000	86,600
AGI: 75th percentile	68,900	103,700	158,900
Observations	15,563,900	63,171,700	36,568,700

Notes: This table presents summary stats for three different samples. The first column uses a random 0.5 percent sample intended to be representative of the population (aged 18 and above) that interacts with the U.S. tax system in a given year, either by appearing on an income tax form as primary, secondary, or dependent filer, or by appearing on an information return. The second column uses the main “ever-IRA” dataset as described in the text (weighted to be representative of that population). The third column restricts to the set of individuals that ever made a Roth conversion between 2003 and 2015. In this table, “married” includes those individuals who are married, filing separately. Quantiles of adjusted gross income, and observation counts, are rounded to the nearest 100 in order to protect taxpayer confidentiality.

Table 2: Cross-tabulation of Decision to Delay Recognition and Optimality of Decision

	(1) Taxpayer recognized immediately	(2) Taxpayer delayed to 2011/2012	(3) Total
0 percent discount rate			
Immediate recognition optimal	0.133	0.191	0.323
Same tax either way	0.102	0.182	0.284
Delaying to 2011/2012 optimal	0.098	0.295	0.392
Total	0.332	0.668	1.000
6 percent discount rate			
Immediate recognition optimal	0.099	0.103	0.202
Same tax either way	0.021	0.006	0.027
Delaying to 2011/2012 optimal	0.213	0.558	0.771
Total	0.332	0.668	1.000

Notes: This table provides two cross-tabulations of (1) the decision by 2010 converters whether to recognize the Roth conversion income immediately or over equal installments in 2011 and 2012 (in the two columns) against (2) whether, according to our calculator, it was optimal to recognize immediately or in 2011 and 2012 (in the three rows). The table additionally provides the marginal distributions in bold. The top panel assumes a zero percent nominal discount rate, while the bottom panel assumes a 6 percent nominal discount rate. The sample is restricted to the set of 2010 converters who file Form 8606 as required and who report “valid” amounts in the applicable fields. See the text for further discussion of these sample restrictions and for details on our tax calculator.

Table 3: Summary statistics: Converters vs. unweighted control

	2004-2009		2010		2011-2015	
	(1) Control	(2) Converter	(3) Control	(4) Converter	(5) Control	(6) Converter
Male (share)	0.494	0.547	0.494	0.544	0.495	0.545
Married (share)	0.669	0.640	0.671	0.773	0.674	0.750
Age (mean)	50.3	51.0	53.3	52.2	56.2	51.2
Age 65+ (share)	0.198	0.234	0.253	0.209	0.317	0.215
IRA balance (median)	5,000	36,400	6,100	50,100	7,800	55,700
Investment income (median)	200	1,100	100	2,500	0	1,900
Workplace Q.P. saving (mean)	1,838	1,717	1,745	4,519	1,740	5,201
Wage share of income (mean)	0.692	0.570	0.660	0.612	0.635	0.621
Any business income (share)	0.245	0.280	0.235	0.372	0.242	0.355
Has Roth IRA	0.172	0.651	0.191	0.554	0.208	0.776
Any Roth Q.P. saving (share)	6	22	27	319	60	638
Income: 25th percentile	30,300	30,900	27,500	55,500	25,300	51,000
Income: median	60,100	54,800	57,200	112,200	55,700	123,900
Income: 75th percentile	104,600	83,500	102,100	217,300	102,300	269,700
Observations	12,594,300	1,095,000	2,109,600	953,100	10,651,500	2,274,900

Notes: This table presents statistics for individuals who make a Roth conversion at time t relative to the remainder of the ever-IRA sample. Income is defined as adjusted gross income, plus non-taxable Social Security, less capital gains, plus IRA deductions and (non-Roth) workplace retirement saving. In this table, individuals who file as married-filing-separately are coded as not married. “Any business income” is a dummy for positive Schedule C or Schedule E income. “Q.P.” saving refers to qualified plan saving, such as in a 401(k). Quantiles of income, and observation counts, are rounded to the nearest 100 to protect taxpayer confidentiality. Dollars amounts are PCE-deflated to 2009 dollars. All variables (except age) are measured as of time $t - 1$.

Table 4: Summary statistics: Converters vs. matched control

	2004-2009		2010		2011-2015	
	(1) Control	(2) Converter	(3) Control	(4) Converter	(5) Control	(6) Converter
Male (share)	0.491	0.545	0.504	0.542	0.502	0.543
Married (share)	0.650	0.641	0.778	0.776	0.758	0.756
Age (mean)	50.9	50.9	52.1	52.1	51.1	51.1
Age 65+ (share)	0.236	0.236	0.208	0.208	0.213	0.213
IRA balance (median)	5,300	35,400	14,800	49,300	19,600	55,000
Investment income (median)	200	1,100	500	2,500	300	1,900
Workplace Q.P. saving (mean)	1,373	1,765	3,614	4,627	4,023	5,295
Wage share of income (mean)	0.648	0.571	0.668	0.614	0.654	0.623
Any business income (share)	0.237	0.283	0.331	0.377	0.351	0.360
Has Roth IRA	0.180	0.653	0.220	0.557	0.270	0.778
Any Roth Q.P. saving (share)	5	22	72	326	186	644
Income: 25th percentile	33,900	33,900	60,000	60,000	55,300	55,300
Income: median	56,900	56,900	116,400	116,400	129,400	129,400
Income: 75th percentile	85,000	85,000	222,000	222,000	275,000	275,000
Observations	1,046,800	1,046,800	920,100	920,100	2,188,200	2,188,200

Notes: This table presents statistics for individuals who make a Roth conversion at time t relative to the matched control group using the coarsened-exact-matching (CEM) weights. Prior to the application of CEM, nonfilers and taxpayers with negative income are dropped. Some additional Roth converters are dropped due to failing the overlap condition required by the CEM procedure. See the text for a further discussion of the CEM procedure. All variables (except age) are measured as of time $t - 1$. See also the notes to Table 3.

Table 5: Workplace retirement saving constraints: Roth converters and control group

	Converters vs. control		Taxable vs. backdoor converters	
	(1) Converter	(2) Matched Control	(3) Taxable	(4) Backdoor
Workplace only				
Ineligible	0.290	0.289	0.323	0.221
At maximum	0.223	0.129	0.169	0.487
<i>Constrained</i>	0.513	0.418	0.492	0.708
Including IRAs				
Ineligible	0.142	0.118	0.155	0.162
At maximum	0.235	0.130	0.191	0.453
<i>Constrained</i>	0.377	0.248	0.346	0.615

Notes: This table reports the share of individuals in certain groups that are constrained from increase retirement saving in certain accounts. Columns 1 and 2 compare converters (2004-2015) to the matched control group, using CEM weights as discussed in Section 5. Columns 3 and 4 compare taxable converters (i.e., those for whom basis represents less than 80 percent of the conversion amount) to backdoor converters (the complement); CEM weights are not used in these two columns. The first row reports the share of individuals who do not work for an employer that appears to offer a qualified (DC) retirement plan. The second row reports the unconditional share of individuals whose qualified plan saving is at least 90 percent of the statutory maximum, including catch-up, if applicable. The third and fourth row repeat the same analysis, broadening the set of retirement plans under consideration to include IRAs. See text for how “ineligible” and “at maximum” are defined in these two rows. Throughout this table, we restrict attention to individuals age 55 or under. Columns 3 and 4 are restricted to 2010 onward.

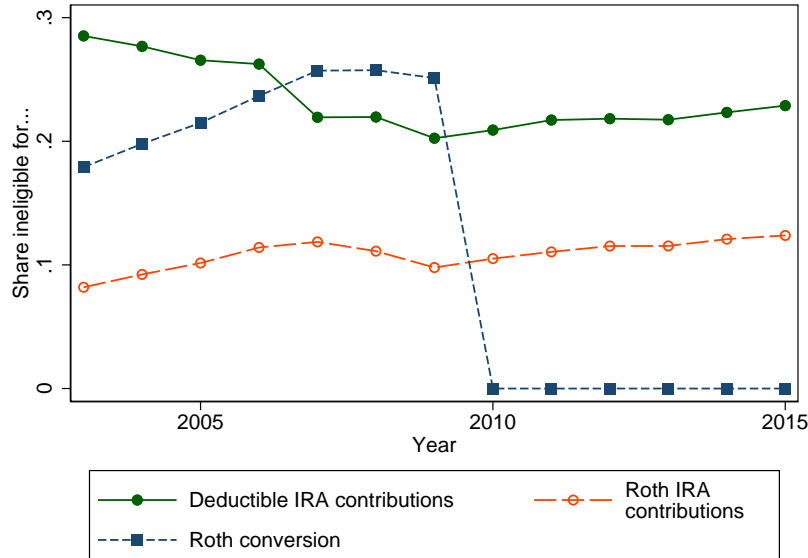
Table 6: Contributions to Roth qualified plans, by converters and control

	(1) Converter	(2) Control	(3) Matched Control
Unconditional	0.071	0.013	
Unconditional (matched converter/control sample)	0.075		0.031
Conditional on availability	0.233	0.076	

Notes: This table presents the share of individuals who save in a Roth option in a workplace qualified plan. The first row compares the converter sample to the rest of the ever-IRA sample. The second row uses CEM weights, as discussed in Section 5. The third row conditions on the individual working for a firm that appears to offer a Roth qualified plan.

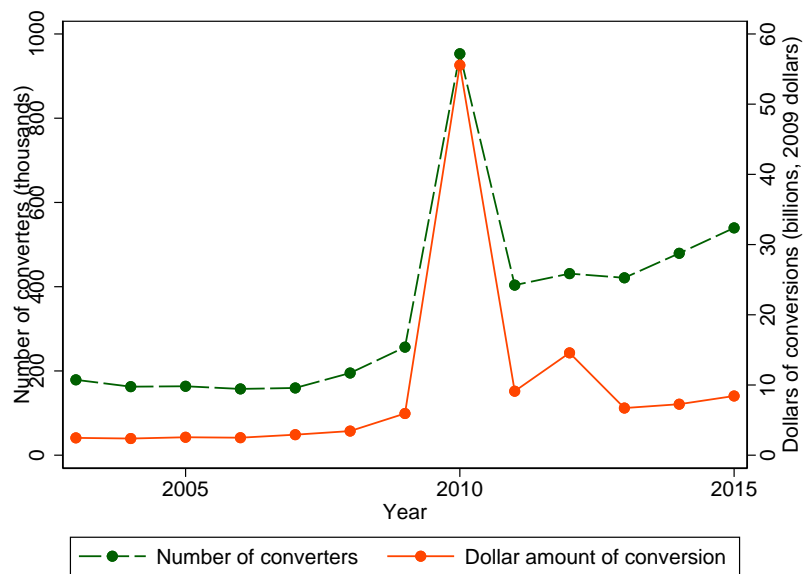
Figures

Figure 1: Share of population ineligible for deductible IRA contributions, Roth IRA contributions, and Roth conversions




Notes: This figure plots the share of individuals that we calculate to have income such that they are ineligible to make a deductible contribution to an IRA, a contribution to a Roth IRA, or a Roth conversion. In determining eligibility to make deductible contribution to an IRA, we additionally use information on the participation of the individual (and his or her spouse, if applicable) in a workplace retirement plan. In constructing these series, we use the “population” dataset, designed to be representative of the U.S. population age 18 and over that interacts with the tax system in a given year.

Figure 2: Number of Roth converters and dollar amount converted



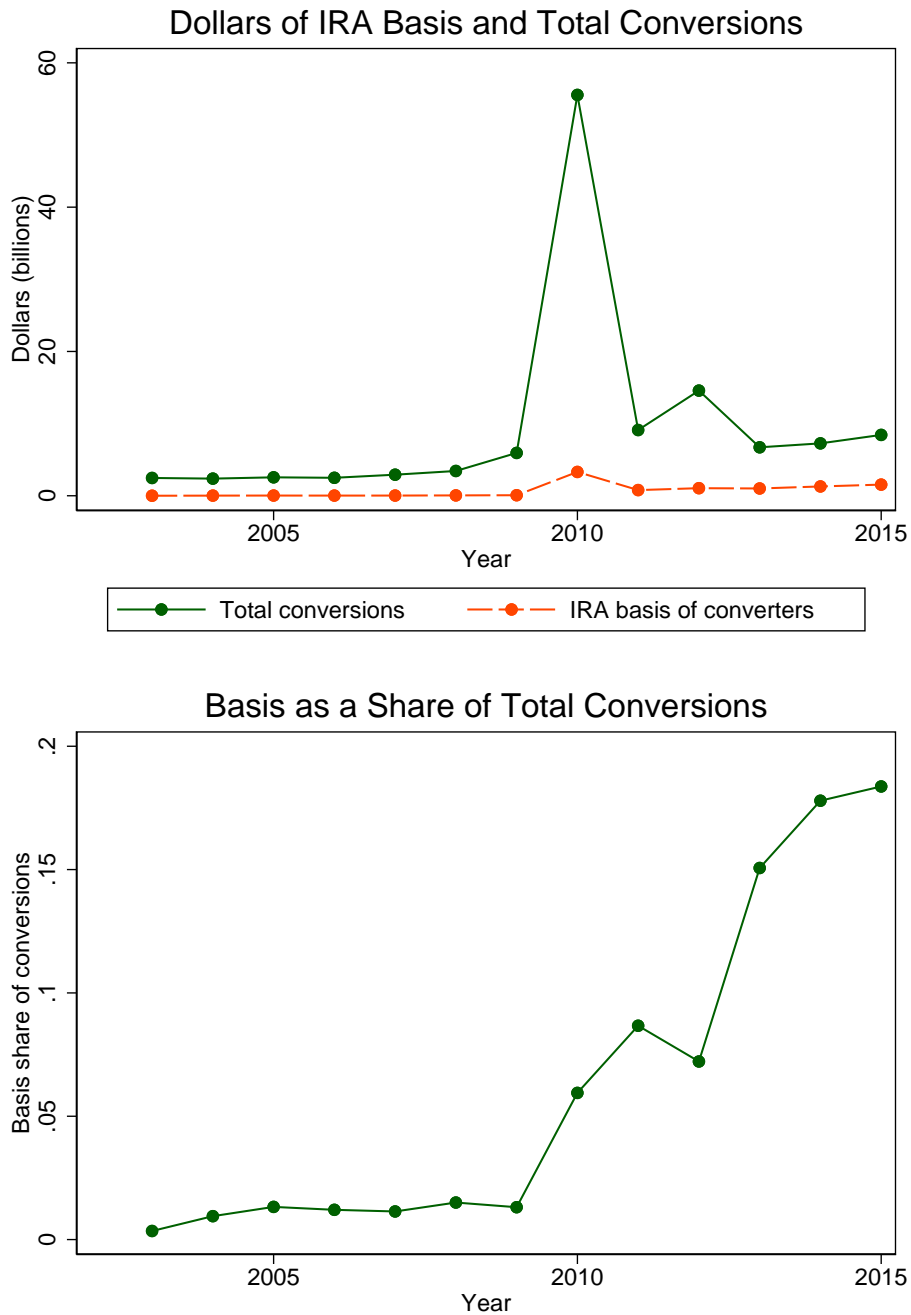
Notes: This figure plots (solid line) the total dollars of Roth conversions that we estimate to have occurred in each year between 2003 and 2015. The dotted line plots the number of individuals who make a Roth conversion in each year. Amounts which are recharacterized back to a traditional IRA are netted out prior to both of these calculations. Dollars amounts are PCE-deflated to 2009 dollars.

Figure 3: Median dollar amount converted



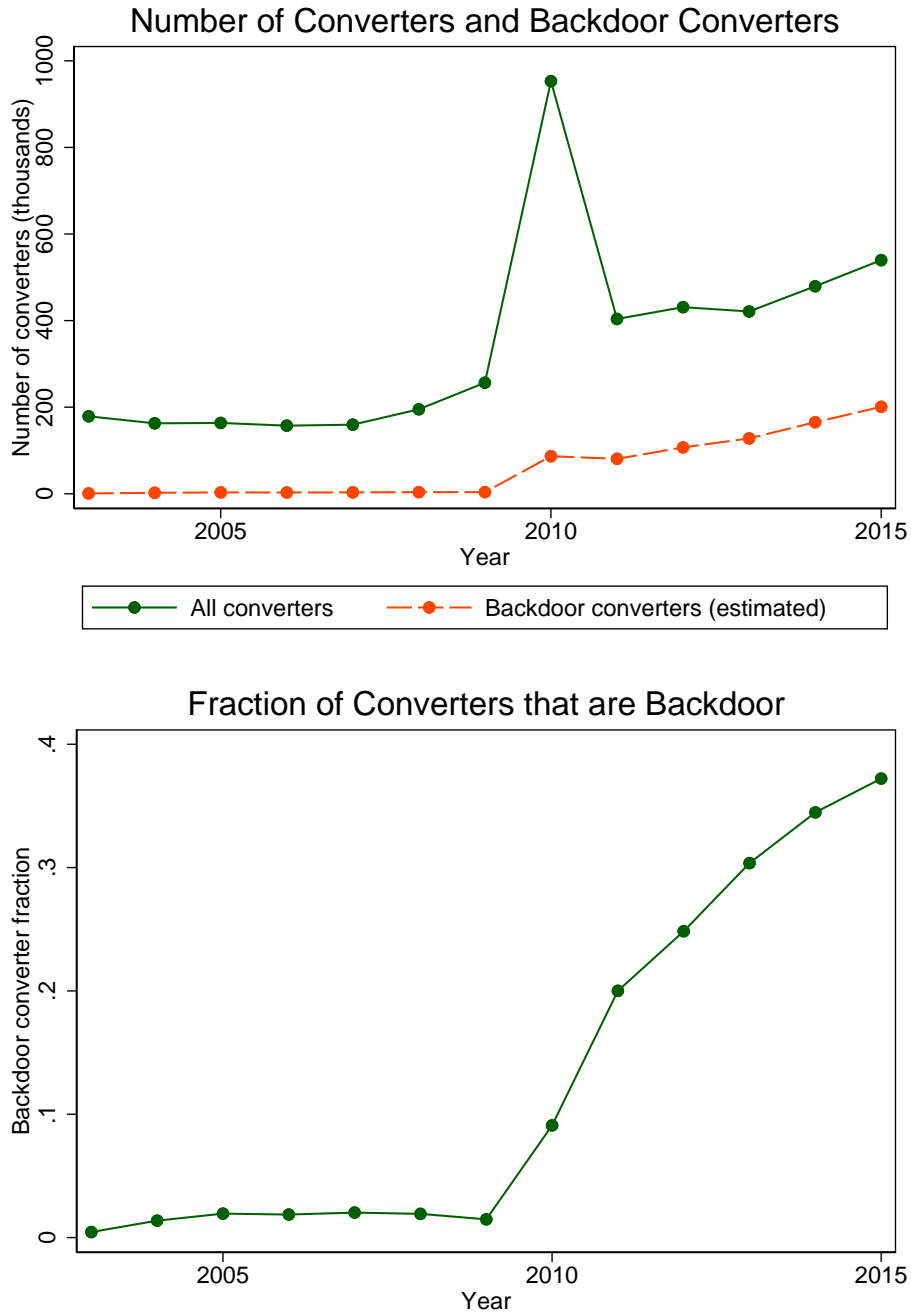
Notes: This figure plots the mean (solid line) and median (dotted line) value of Roth conversions in each year from 2003 to 2015. Dollar amounts are PCE-deflated to 2009 dollars.

Figure 4: Dollars of basis, relative to total conversions



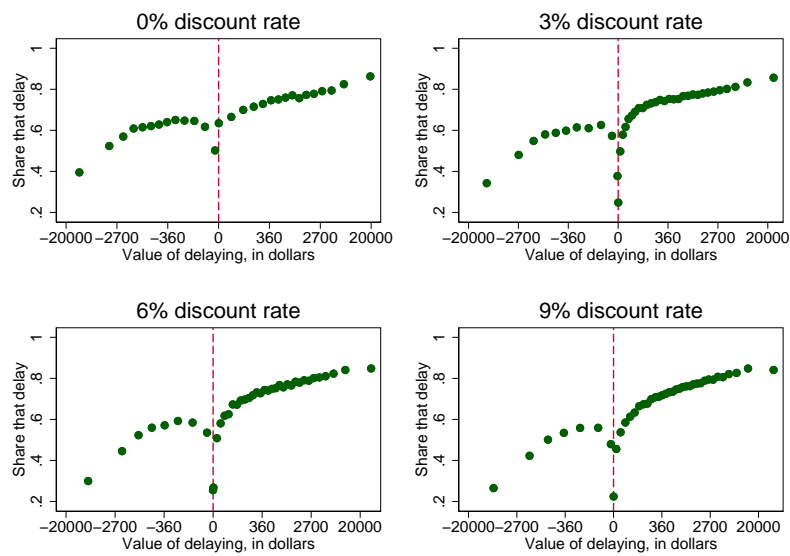
Notes: The top panel of this figure plots the total amount of Roth conversions and the total amount of basis in those conversions, over time. The bottom panel plots the fraction between these two series. Basis is calculated based on estimated non-deductible IRA contributions and other factors, as discussed in Section 3. The figures are created using all converters. Dollars are PCE-deflated to 2009 dollars.

Figure 5: Number of backdoor converters, relative to total converters



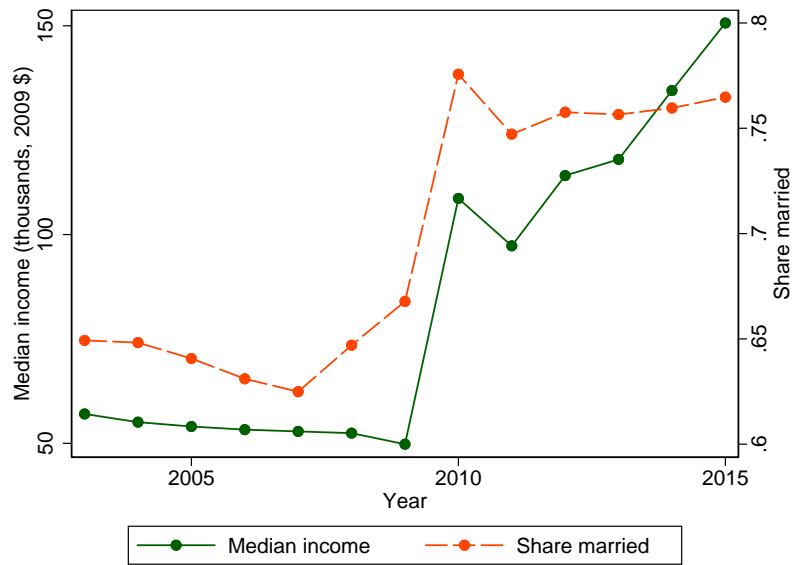
*Notes:*The top panel of this figure plots the total number of converters (solid line) and the the number of backdoor converters (dotted line), over time. The bottom panel plots the ratio of these two series. Backdoor conversions are defined as having basis at least 80% of total conversion. Basis is calculated based on estimated non-deductible IRA contributions and other factors, as discussed in Section 3. This uses all converters.

Figure 6: Delaying recognition to 2011/2012 as a function of the benefit of delaying



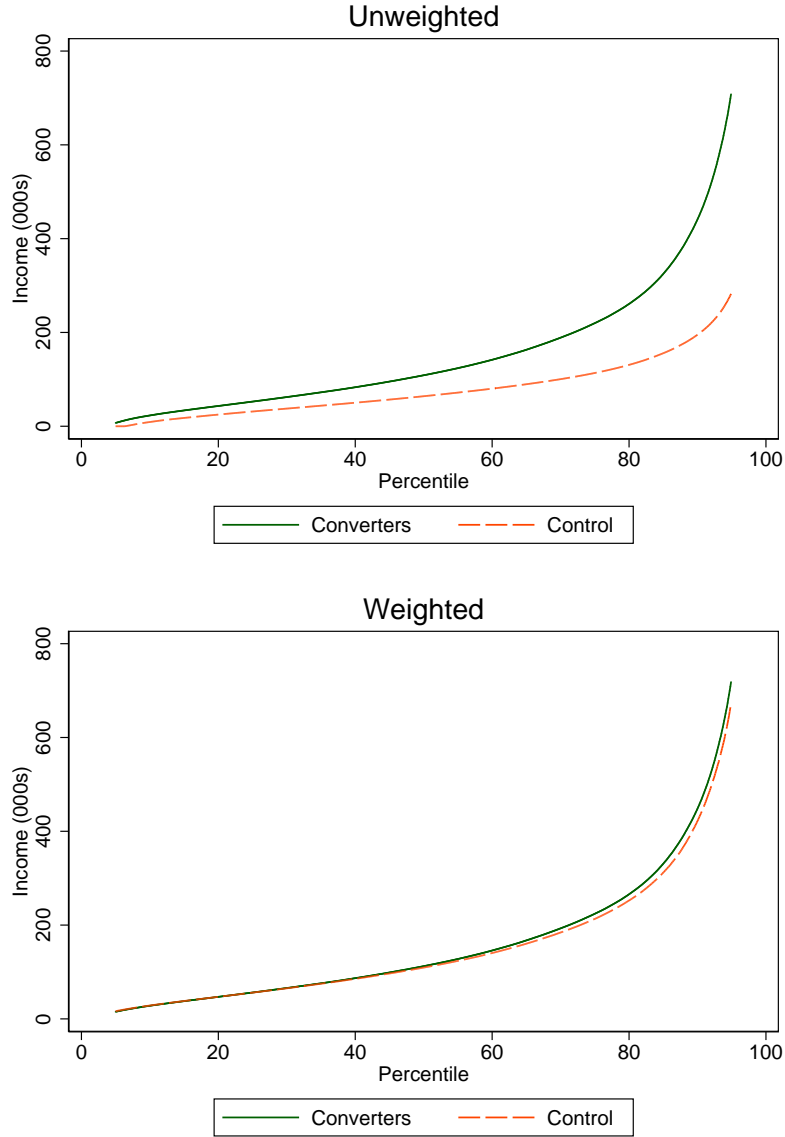
Notes: The four panels of this figure plot binned scatter plots of the share of 2010 converters who choose to delay recognition of Roth conversion income until 2011 and 2012, as a function of the tax benefit of doing so. Each panel assumes a different discount rate for the purpose of calculating this tax benefit. The x-axis is scaled using a modified inverse hyperbolic sine scale. The sample is restricted to those who correctly file Form 8606 with a positive taxable amount, and who make valid entries on the applicable lines, as discussed in the text. See Section 3 for further details about the tax calculator.

Figure 7: Characteristics of converters



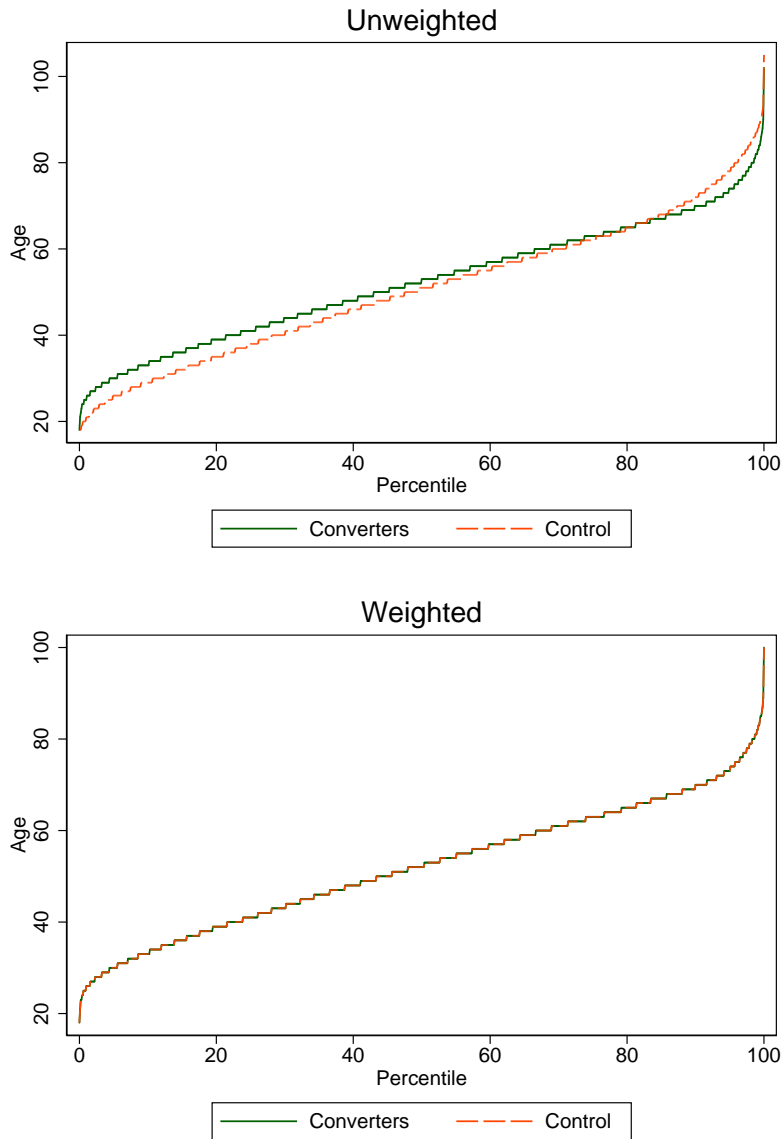
Notes: This figure plots the median income and the share married among the set of Roth converters, over time. Income is defined as adjusted gross income, plus non-taxable Social Security, less capital gains, plus IRA deductions and (non-Roth) workplace retirement saving. Dollar amounts are PCE-deflated to 2009 dollars.

Figure 8: Distribution of income of 2010 converters with and without reweighting




Notes: The top panel of this figure plots the quantile function (the inverse of the cumulative distribution function) for prior year income for 2010 converters, relative to the rest of the ever-IRA sample. The bottom panel re-weights the sample using the CEM weights, as discussed in Section 5. Income is defined as adjusted gross income, plus non-taxable Social Security, less capital gains, plus IRA deductions and (non-Roth) workplace retirement saving. Dollar amounts are PCE-deflated to 2009 dollars. For readability, percentiles below 5 and above 95 are not shown.

Figure 9: Distribution of age of 2010 converters with and without reweighting



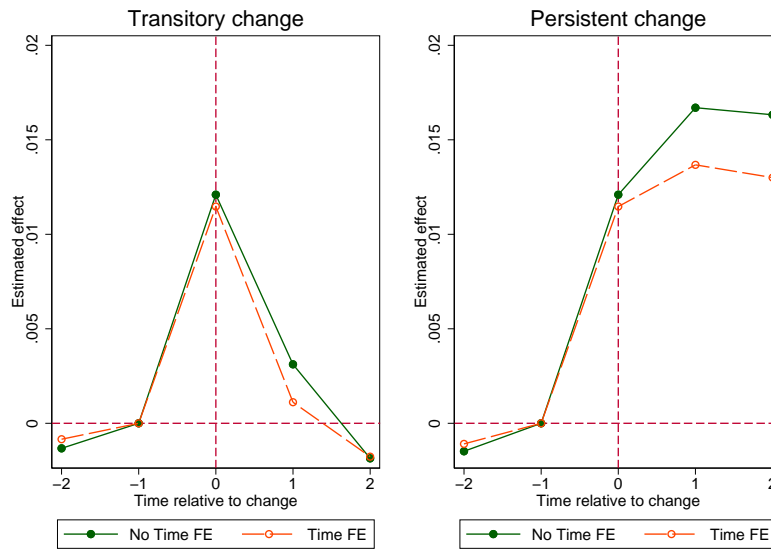
Notes: The top panel of this figure plots the quantile function (the inverse of the cumulative distribution function) for age for 2010 converters, relative to the rest of the ever-IRA sample. The bottom panel re-weights the sample using the CEM weights, as discussed in Section 5.

Figure 10: 2012 and 2011 Roth conversions, by income group




Notes: This panel plots aggregate Roth contributions in 2011 (dotted line) and 2012 (solid line) among different income groups. Income groups are constructed such that the 2011 Roth conversions are the same within each group, which causes the 2011 series to be flat, by construction. Individuals are assigned to income bins based on their income in that year. Those with income less than zero are excluded. Income is defined as adjusted gross income, plus non-taxable Social Security, less capital gains, plus IRA deductions and (non-Roth) workplace retirement saving. Dollar amounts are PCE-deflated to 2009 dollars.

Figure 11: Effect of changes in MTR on probability of Roth conversions




Notes: The left panel in this figure plots $\beta_k = \tilde{\beta}_k - \tilde{\beta}_{-1}$ from the following regression, where y_{it} is a dummy for making a Roth conversion and mtr_{it} is the marginal tax rate on the first dollar of IRA distributions, scaled as a number (typically) between 0 and 1: $\Delta y_{it} = \delta_t + \sum_{k=-2}^2 \tilde{\beta}_k \Delta(1 - mtr_{i,t-k}) + \epsilon_{it}$. These coefficients represent the effect of a transitory decrease in marginal tax rates. The right panel plots the cumulated coefficients, relative to the cumulated coefficient for $k = -1$. See text for how these cumulated effects are calculated. These represent the implied effect (relative to $k = -1$) of a persistent increase in marginal tax rates. In both panels, the series with the solid line and solid circles omits the time fixed effects.

Figure 12: “Complete” Roth Conversions



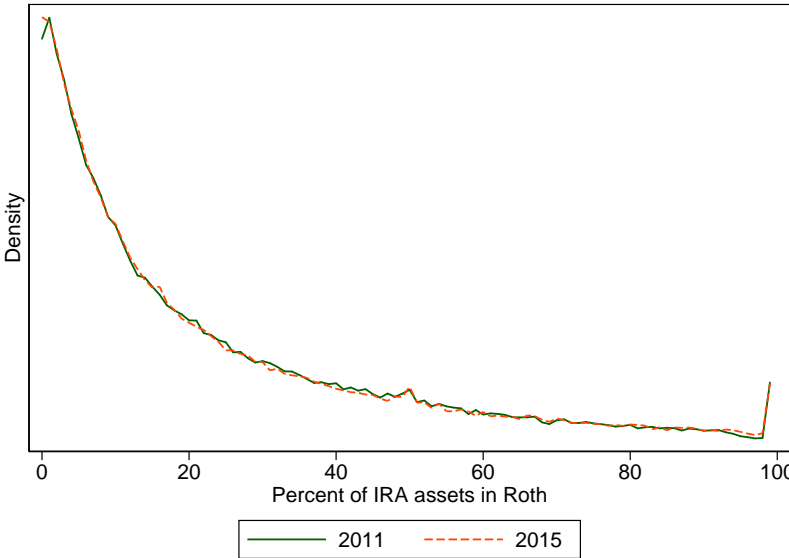
Notes: This figure plots the share of non-backdoor Roth conversions that are “complete”. We define a Roth conversion to be “complete” if total traditional IRA assets are zero at the end of the year. For this purpose, traditional IRA assets do not include SEP and SIMPLE accounts. Backdoor conversions are defined as having basis share at least 80 percent.

Figure 13: Share of non-complete 2010 converters with zero IRA assets, 2011-2015




Notes: This figure plots the share of 2010 non-complete, non-backdoor converters that have zero traditional IRA assets in subsequent years. For this purpose, traditional IRA assets do not include SEP and SIMPLE accounts. Backdoor conversions are defined as having basis share at least 80 percent.

Figure 14: Density of Roth share of IRA assets, among those with positive traditional and Roth assets, in 2011 and 2015



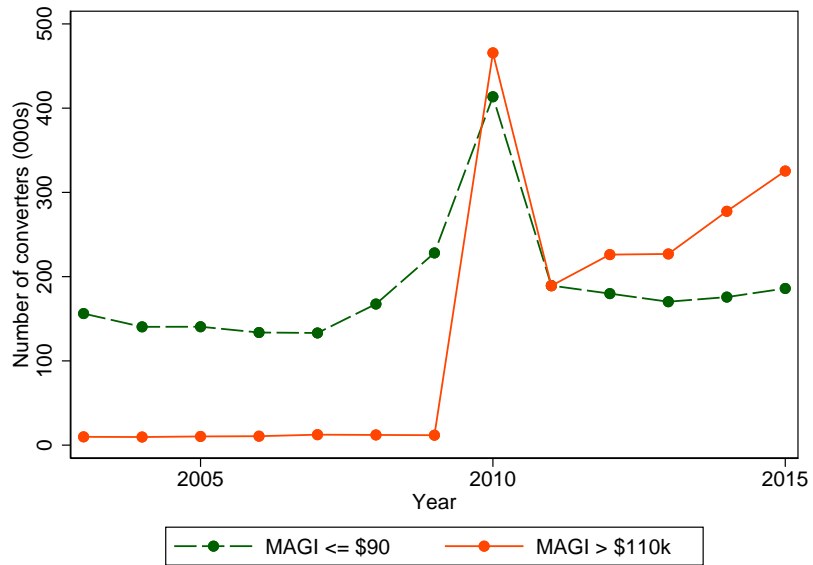
Notes: This figure plots the ratio of Roth IRA assets to (Roth IRA assets plus traditional IRA assets), among the subset of individuals with both positive Roth IRA assets and traditional IRA assets, separately in 2011 and 2015. For this purpose, traditional IRA assets do not include SEP and SIMPLE accounts.

Figure 15: Are individuals converting in order to gain liquidity? Roth IRA distributions, by year relative to conversion




Notes: This figure plots the results of the following event-study regression, where τ indexes the year of the Roth conversion, t indexes calendar years, and $y_{it\tau}$ is a dummy for making any Roth IRA distributions (with distribution codes J, Q, or T): $y_{it\tau} = \alpha_t + \sum_{k=-2, k \neq -1}^5 \beta_k \times \mathbb{1}(t - \tau = k) + \epsilon_{it\tau}$. The regression is estimated on an eight-year panel (corresponding to event-times -2 through 5) surrounding each person-by-conversion ($i\tau$) observation. If an individual converts multiple times during this period, they enter the data as multiple different panels. These panels are balanced, except for death and it observations where the individual is younger than 18. Event-time $k = -1$ is omitted from the regression; the corresponding value for $k = -1$ in the figure is the raw share of converters making a Roth IRA distribution at time -1 . Other values are equal to the coefficient from the regression, plus this $k = -1$ baseline. At event time $k = 5$, individuals may begin taking distributions from the Roth IRA, up to amount of the conversion, without penalty. Standard errors are clustered by individual. We restrict attention to individuals who are no older than 53 at the time of conversion, who therefore are younger than the 59.5 threshold throughout the entire panel.

Figure 16: Number of converters below and above \$100,000 of income



Notes: This figure plots the number of individuals making a Roth conversion by year in two bins of modified adjusted gross income (MAGI): below \$90,000, and above \$110,000, in nominal dollars. Prior to 2010, individuals with MAGI in excess of \$100,000 were not allowed to perform a Roth conversion.

Figure 17: Effect of RMDs: Probability of conversion around age 70.5



Notes: This figure plots the fraction of individuals making a Roth conversion, conditional on having a positive IRA balance at the end of the prior year, by exact date of birth in seven-day bins. The sample is restricted to those turning age 70 in the year in question. Individuals to the left of zero turn 70.5 prior to the end of the year, and thus are required to take an RMD for that year. Individuals to the right of zero do not turn 70.5 prior to the end of the year, and thus are not required to take an RMD for that year. The data is pooled across all years, except that 2009 is dropped, since RMD rules were suspended in 2009.

Appendix Figures and Tables

Table A1: Cumulative lag/lead coefficient estimates

	Transitory tax change		Persistent tax change	
	No time FE	Time FE	No Time FE	Time FE
$t = -2$	-0.001*** (0.000)	-0.001*** (0.000)	-0.001*** (0.000)	-0.001*** (0.000)
$t = -1$	0 (0)	0 (0)	0 (0)	0 (0)
$t = 0$	0.012*** (0.000)	0.011*** (0.000)	0.012*** (0.000)	0.011*** (0.000)
$t = 1$	0.003*** (0.000)	0.001*** (0.000)	0.017*** (0.000)	0.014*** (0.000)
$t = 2$	-0.002*** (0.000)	-0.002*** (0.000)	0.016*** (0.000)	0.013*** (0.000)
Observations	38,657,000	38,657,000	38,657,000	38,657,000
Time fixed effects		X		X

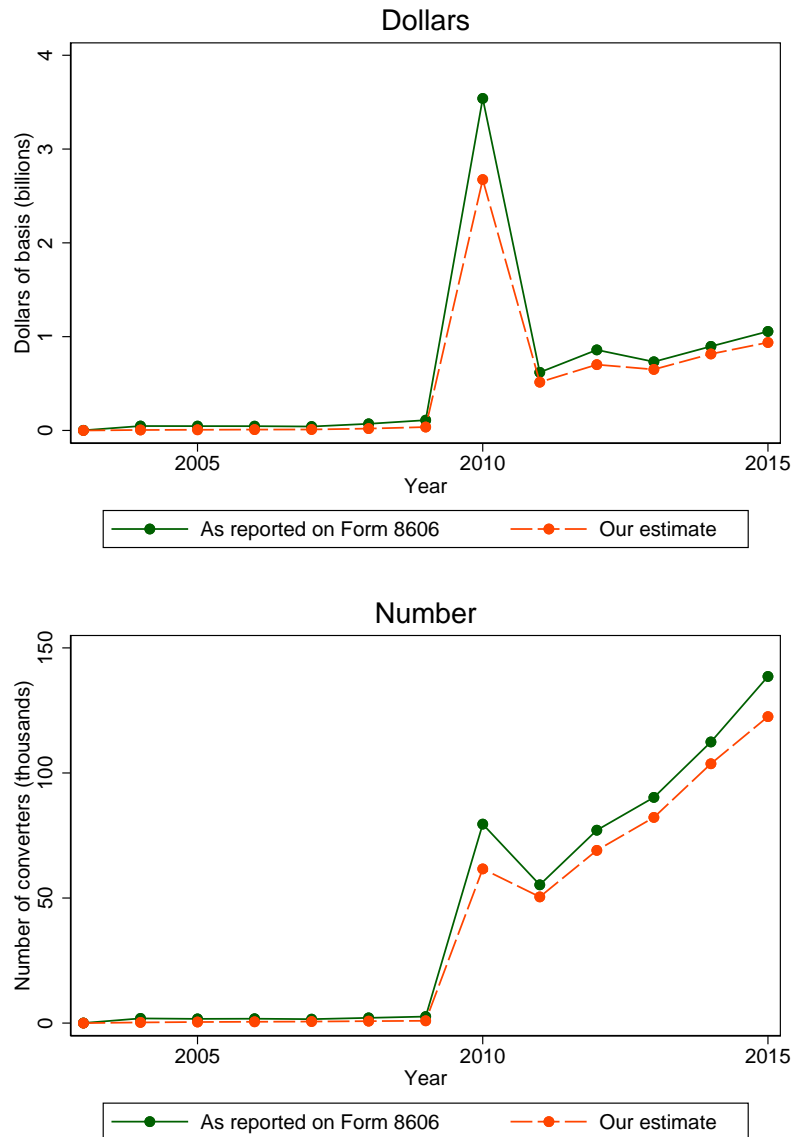
Notes: This table reports the coefficient estimates plotted in Figure 11. The left two columns report $\beta_k = \tilde{\beta}_k - \tilde{\beta}_{-1}$ from the following regression, where y_{it} is a dummy for making a Roth conversion and mtr_{it} is the marginal tax rate on the first dollar of IRA distributions, scaled as a number (typically) between 0 and 1: $\Delta y_{it} = \delta_t + \sum_{k=-2}^2 \tilde{\beta}_k \Delta(1 - mtr_{i,t-k}) + \epsilon_{it}$. These coefficients represent the effect of a transitory decrease in marginal tax rates. Columns 3 and 4 report the cumulated coefficients, relative to the cumulated coefficient for $k = -1$. See text for how these cumulated effects are calculated. These represent the implied effect (relative to $k = -1$) of a persistent increase in marginal tax rates. In both panels, the series with the solid line and solid circles omits the time fixed effects. * $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$.

Table A2: Regression discontinuity estimates of being subject to RMD on making a Roth conversion

	(1)	(2)	(3)	(4)
	BW: 30 day	BW: 60 day	BW: 90 day	BW: 120 day
Subject to RMD	-0.00208*** (0.000436)	-0.00202*** (0.000261)	-0.00193*** (0.000201)	-0.00184*** (0.000167)


Notes: This table reports the regression discontinuity estimates of the effect of being subject to RMDs on the probability of making a Roth conversion. Specifically, these are the coefficient estimates from the following regression, where y_{it} is a dummy for making a Roth conversion in year t , where x_i is age in days relative to 70.5 at the end of year t , and $f(x_i)$ is a continuous one-degree polynomial, whose slope is allowed to differ on each side of 0: $y_{it} = f(x_i) + \beta \times \mathbb{1}(x_i \geq 0) + \epsilon_{it}$. The sample is restricted to those with positive IRA balance at time $t - 1$. The regression is estimated using local linear regression with bandwidths specified in the column headings.

Figure A1: Sanity check: How well does estimated basis match reported basis among Form 8606 filers?




Notes: The top panel of this figure plots the total basis reported by taxpayers on Form 8606 (solid line) and our calculated basis, among the subset of Roth converters that file Form 8606 (dashed line). The bottom panel reports the number of backdoor converters (defined as having basis in excess of 80 percent of the conversion) among those who file Form 8606; the solid line uses basis as reported by taxpayers while the dashed line uses our calculated basis. Dollar amounts are PCE-deflated to 2009 dollars.

Figure A2: Delaying recognition to 2011/2012 as a function of conversion size




Notes: This figure plots a binned scatter plot of the share of 2010 converters that delay recognition of Roth conversion income until 2011 and 2012, as a function of the size of the conversion. This figure makes the restrictions discussed in Section 4.3.

Figure A3: Testing the RD assumption: distribution of age (in days) relative to 70.5



Notes: This figure plots the density of observations with respect to exact date of birth (in seven-day bins) for the sample included in 17. The sample is restricted to those turning age 70 in the year in question. Individuals to the left of zero turn 70.5 prior to the end of the year, and thus are required to take an RMD for that year. Individuals to the right of zero do not turn 70.5 prior to the end of the year, and thus are not required to take an RMD for that year. The data is pooled across all years, except that 2009 is dropped, since RMD rules were suspended in 2009.

Figure A4: Placebo regression discontinuity: Probability of making a Roth conversion around age 68.5



Notes: This figure plots the fraction of individuals making a Roth conversion, conditional on having a positive IRA balance at the end of the prior year, by exact date of birth in seven-day bins. The sample is restricted to those turning age 68 in the year in question. All individuals in this figure are *not* required to take an RMD for this year. The data is pooled across all years, except that 2009 is dropped, since RMD rules were suspended in 2009.