PRELIMINARY

The Employment Effects of the Social Security Earnings Test

Alexander Gelber
UC Berkeley Goldman School of Public Policy and NBER

Damon Jones
University of Chicago Harris School of Public Policy and NBER

Daniel Sacks
Indiana University Kelley School of Business

Jae Song
Social Security Administration

August 2017

Abstract

We investigate the impact of the Social Security Annual Earnings Test (AET) on the employment decisions of older Americans. The AET reduces Social Security benefits by one dollar for every two dollars earned above the exempt amount. Using a differences-in-differences design, we compare employment rates subsequent to reaching the Social Security retirement age of those previously earning above and below the AET exempt amount, who form the "treatment" and "control" groups, respectively. We find that the employment rate of those subject to the AET decreases significantly relative to those not subject to it. The point estimate suggests that the AET reduces the employment rate of older Americans aged 63-64 by at least 3.7 percentage points.

This research was supported by the U.S. Social Security Administration through grant #RRC08098400-09 to the National Bureau of Economic Research as part of the SSA Retirement Research Consortium. The findings and conclusions expressed are solely those of the author(s) and do not represent the views of SSA, any agency of the Federal Government, or the NBER.

I. Introduction

The Social Security Annual Earnings Test (AET) has a large effect on Social Security Old Age and Survivor Insurance (OASI) benefits, and therefore could have an important effect on the employment rate of older workers. The AET reduces OASI claimants' current OASI benefits as a proportion of earnings, once a claimant earns in excess of an exempt amount. For example, for OASI claimants aged 62 to 65 in 2016, current OASI benefits are reduced by 50 cents for every extra dollar earned above $15,720. Our main hypothesis is that the AET may reduce OASI claimants' incentives for additional work and lead them to earn less or even stop working.

We investigate the effects of the AET on the employment rate of older workers who are eligible to begin claiming OASI benefits. We propose a new strategy for investigating the effects of the AET on the employment rate by examining how employment outcomes vary with the change in the incentive to work that occurs when one earns above the AET exempt amount.

Our results promise to be relevant to current policy debates regarding the employment rate of older workers. The AET's very high benefit reduction rates raise the possibility that the AET could have large effects on earnings and employment. Indeed, the AET has an important effect on current AET benefits. We estimate using the Social Security Administration Benefits and Earnings Public Use File that in the latest data available from 2003, the AET reduced current OASI benefits for over half a million (around 538,000) individuals, whose current OASI benefits on average were reduced by over half (51.4 percent)—amounting to $4.3 billion dollars' worth of reductions in current benefits.[1] The importance of the AET is reinforced by the fact that the

[1] If earnings and/or employment are reduced by the AET, then the calculated reduction in benefits reflects a lower bound on the reduction in current benefits we would hypothetically observe if earnings and employment were inert in response to the AET. Thus, this strengthens our case that the AET substantially affects current OASI benefits. As

share of older workers affected by this policy—those under the Normal Retirement Age (NRA)—is increasing over time, as the NRA gradually increases from 65 to 67.

Most previous literature on the AET has examined how the AET affects the choice of how much to earn, given that an individual earns a positive amount (e.g. Burtless and Moffitt 1985; Friedberg 1998, 2000; Song and Manchester 2007; Gelber, Jones, and Sacks 2013). For example, the AET could reduce the incentive for additional earnings above the exempt amount and cause an individual to choose part-time work rather than full-time work. However, in principle one of the most important potential effects of the AET could be on the decision of whether or not to work. In other words, because the AET penalizes work by reducing current OASI benefits, some individuals could choose not to work at all. Using a differences-in-differences empirical strategy, some literature has found no evidence that the AET affects the choice to work at all (Gruber and Orszag 2003; Song and Manchester 2007; Haider and Loughran 2008), while other literature has found some evidence for an effect (Friedberg and Webb 2009).

We revisit this conclusion using a different and novel methodological approach. In particular, we focus on employment patterns among those with earnings above and below the exempt amount where the AET begins to take effect, e.g. earnings near $15,720 in 2016. At this exempt amount the incentives to work change notably due to the AET, as the imposition of the AET above the kink point could cause some individuals to exit employment. Using a differences-in-differences design, we compare employment rates subsequent to reaching the Social Security retirement age of those previously earning above and below the AET exempt

discussed below, when current OASI benefits are reduced due to the AET, future benefits may be enhanced; however, our point here is simply that the AET greatly affects benefits, at least in their timing, and therefore has important potential implications.

amount, who form the "treatment" and "control" groups, respectively. Our results show much larger effects on employment than most previous literature had indicated. We find that the employment rate of those subject to the AET decreases significantly relative to those not subject to it. Our point estimates suggest that the AET reduces the employment rate of older Americans aged 63-64 by at least 3.7 percentage points.

Recent work by Gelber, Jones, Sacks, and Song (2017), which also studies workers below and above the exempt amount, represents an additional exception to the conclusion that the AET has little effect on older workers' employment rate. Gelber, Jones, Sacks, and Song (2017) use a localized approach to demonstrate that the AET reduces employment for workers who earn just above the AET exempt amount, as we describe in detail below. In the current study, we expand analysis to a larger set of workers and use an alternative method, therefore reinforcing and extending the conclusions of Gelber, Jones, Sacks, and Song (2017) with a complementary method, which relies on alternative identification assumptions.[2]

The paper proceeds as follows. Section II describes the policy environment. Section III explains the empirical strategy. Section IV describes our data. Section V proceeds to the results. Section VI concludes.

II. Policy environment

OASI provides annuity income to older Americans and to survivors of deceased workers. Individuals with sufficient years of eligible earnings can claim OASI benefits through their own work history as early as age 62, the Early Entitlement Age (EEA). They can claim full benefits once they reach the Normal Retirement Age (NRA), which is 65 for individuals in our sample. The AET reduces current OASI benefits in proportion to earnings above an exempt amount

[2] Because the policy environment and data are similar to Gelber, Jones, Sacks, and Song (2017), these sections of this paper have overlap with this previous paper.

shown in Figure 1 over the years we study. For those whose age is at the NRA and over, the real exempt amount is substantially higher than for those below the NRA. For those under the NRA but above the EEA—the main group that our empirical work studies—the benefit reduction rate (BRR) was 50 percent throughout the period we study, 1978 to 1987. During our period, the AET applied to earnings from ages 62 to 71 during 1978 to 1982, and from ages 62 to 69 during 1983 to 1987.

The AET applies only to an individual's earnings; spouses' earnings do not count in the earnings total to which the AET is applied. For a retired worker (i.e. primary) beneficiary whose spouse collects spousal benefits, the AET reduces the family's total OASI benefit by the amounts we have described. The family's total benefit is also reduced at the benefit reduction rate when the spouse (separately) earns more than the AET threshold. For a retired worker beneficiary whose spouse is collecting benefits on his or her own earnings record, the AET reduces the retired worker beneficiary's benefits by the amounts described while not affecting the spouse's benefits.

When current OASI benefits are lost to the AET, future scheduled benefits may be increased in some circumstances. This is sometimes referred to as "benefit enhancement." For beneficiaries below the NRA in particular, the benefit enhancement, known as the "actuarial adjustment," raises future benefits whenever a claimant earns over the AET exempt amount. Future benefits are raised by 0.55 percent per month of benefits withheld for the first three years of AET assessment. On average, this actuarial adjustment is considered roughly actuarially fair (Diamond and Gruber 1999).

Despite the existence of benefit enhancement, individuals could still perceive the AET as penalizing current earnings, for several potential reasons. First, the AET benefit enhancement

was on average roughly actuarially fair only beginning in the late 1990s. Second, for liquidity-constrained individuals, those whose expected lifespan is shorter than average, or those who discount particularly quickly, the AET is more punitive—and such individuals could also choose to reduce work or stop participating in the labor force in response to the AET. Finally, many individuals also may not understand the AET benefit enhancement or other aspects of OASI (Liebman and Luttmer 2015; Brown, Kapteyn, Mitchell, and Mattox 2013). Previous literature has found significant bunching responses to the AET (e.g. Friedberg 2000; Gelber, Jones, and Sacks 2013), implying that some individuals act as if the AET is punitive.

III. Empirical strategy

Our approach involves a difference-in-difference estimator of the effect of the AET on employment. Specifically, we will exploit the fact that the AET only affects those who would earn above the exempt amount in a given year if they choose to work. Furthermore, the AET only affects those who claim Social Security, which can be claimed as early as age 62. We therefore will compare those who would potentially earn above and below the exempt amount, before and after turning 62.[3]

Our analysis must take account of the fact that one's current earnings may also respond to the AET, and, furthermore, we do not observe earnings for those who exit the labor force. In order to circumvent this problem, we will use lagged earnings as a proxy for potential earnings in the present period. To support the validity of using lagged earnings as a proxy for future desired earnings, we show that desired earnings remain stable across a "placebo" set of ages. Specifically, we show that the distribution of real earnings growth from one period to a

[3] Due to the nature of the data, we cannot consistently estimate a husband's response to a wife's incentives or vice versa. We only observe husbands linked to their wives when one spouse is collecting as a dual or secondary beneficiary, which is a highly selected sample.

subsequent period exhibits a spike at zero. Figure 2 shows that from age 59 to age 60—a placebo set of ages during which our sample is not subject to the AET—a noticeable spike in real earnings growth does occur near zero percent growth.

Thus, given positive earnings z_{ai}, for individual i at "base" age a, we will examine the probability that one continues to have positive earnings at age $a + t$, where t =3, 4 or 5. We explain below why we consider a minimum of lead of 3. We denote the AET exempt amount by z^* and compare the difference in the probability of working for those with $z_{ai} \geq z^*$ to those with $z_{ai} < z^*$ when $a + t < 63$ and when $a + t \geq 63$. We treat age 63 as the post period because that is the first year we could see an extensive margin effect of the AET.[4]

Thus, our approach constitutes a "differences-in-differences" analysis. Our outcome is the probability of having zero earnings t years in the future. Our "treatment" group is individuals with earnings above the exempt amount at age a, and our "control" group is individuals with earnings below the exempt amount at age a. Our pre-period (i.e. prior to the potential effect of the AET) is when $a + t < 63$, and our post-period (i.e. after the potential effect of the AET) is when $a + t \geq 63$.[5]

In regression form, our difference-in-difference model will be estimated using the following specification:

$$E_{i,a+t} = \alpha_0 + \gamma_a + \delta \cdot 1\{z_{ia} \geq z^*\} + \beta \cdot 1\{z_{ia} \geq z^*, a + t \geq 63\} + \varepsilon_{ia}, \qquad (1)$$

[4] The AET first applies to claimants when they reach OASI eligibility at age 62, but it does not make sense to examine the effect of the AET on whether an individual has positive earnings in the calendar year s/he turns age 62. The reason is that we observe calendar year earnings. If an individual claims OASI at age 62, the AET only applies to earnings in the months after the individual claims. If the claimant earns at all during this calendar year—even during months prior to claiming OASI—then she will have positive earnings in this calendar year. Thus, a person who is induced by the AET to stop earning after claiming would appear in the data with positive earnings during this calendar year, and therefore would appear to have no measured response to the AET.

[5] Our strategy implicitly assumes that desired earnings change modestly from age 60 to age 63. Indeed, our preliminary results show that for a large percentage of workers in a placebo sample three years apart (ages 57 and 60), there is little change in real earnings.

where $E_{i,a+t}$ is a binary variable that equals one if an individual is employed t years after age a.

We can estimate our main specification in equation (1) using ordinary least squares (OLS). We allow t to vary between three, four, and five years as alternative robustness checks. The key independent variables are a function of earnings in the baseline age a for individual i, i.e. $z_{i,a}$. The parameter γ_a is an age-specific fixed effect and δ is a time-invariant, level difference in the probability of having positive earnings. The key parameter of interest is β, the average difference in the probability of having positive earnings at age 63 or older, for those who would earn above the exempt amount, z^*, relative to those who would earn below the exempt amount.

Our method relies on two key assumptions. First, we make the "parallel trends" assumption, common to all difference-in-difference estimators. That is, we assume that in the absence of our key policy variation, i.e. the introduction of the AET, the difference in the probability of positive earnings between those with earnings below and above the exempt amount would only differ by a constant amount. Formally, we impose the restriction in equation (1) that the parameter δ is time-invariant. We cannot technically test this assumption directly, but we can assess the merits of this assumption using data from ages younger than 63. Our assumption implies that labor force participation during these ages follows a similar trend for those earnings below versus above the exempt amount.

Second, we assume that earnings at age a are a reasonable proxy for potential earnings at age $a + t$, conditional on working and in the counterfactual scenario where the AET is not in effect. We would ideally model the probability positive earnings at age $a + t$ as a function of whether or not an individual is affected by the AET in that same year. However, we cannot observe potential earnings at age $a + t$ for those who do not work, and for those who do work,

earnings may be decreased in response to the AET. Instead, we will use a lagged measure of earnings from age a, as a proxy for the earnings that would be realized at age $a + t$. This assumption will be violated if individuals tend to adjust their earnings in anticipation of the AET. However, our prior research has indicated that during the ages used in our analysis, there is little evidence of anticipatory adjustment to the AET (Gelber *et al.*, 2013 2017). In order to further assess the validity of this assumption, we can gauge how earnings evolve over a window of one to three years, during ages prior to age 63. Our investigation has revealed that subsequent earnings are highly predictable over this time frame.

These considerations also dictate our choice of using a 3-year gap between base age earnings and the employment outcome we investigate. We have to use at least a two year lag to look at the age 63 effect of the AET (because age 62 earnings could respond to the AET).[6] Our differences-in-differences method will control for the possibility that high earners have higher labor force attachment by removing the difference attributable to the constant effect of being in the low- or high-earning group at age *a*. The method relies on the assumption that in the absence of the AET, the difference in the probability of employment would be constant between those earning above and below the exempt amount, before and after age 63. While we cannot directly test this assumption, we can perform basic checks on its validity.

First, a testable prediction is that we should see the difference in employment probabilities between these two groups evolving similarly prior to age 63. Second, in our differences-in-differences specification, we can add separate trends in age in the below-z^* and

[6] We use a three-year lag to guard against the possibility of anticipatory adjustment in age 61 earnings to the expected future imposition of the AET (and demonstrate that there is no evidence of anticipatory adjustment at age 60).

above-z^* groups, to demonstrate that a significant difference occurs precisely around the time when the AET is imposed, relative to this trend.

If those with earnings below the AET exempt amount have weaker labor force attachment and hypothetically drop out of the labor force more when they reach retirement age, then we would expect a larger decrease in the employment probability at retirement age among those with earnings below the exempt amount than among those above the exempt amount. This is the opposite of what we find in our results, namely a *smaller* decrease in the employment probability at retirement age among those with earnings below the exempt amount than among those above the exempt amount. Thus, if anything this hypothesis would push against our finding of a large response.

IV. Data

We implement our estimation strategy using the restricted-access Social Security Administration Master Earnings File (MEF) linked to the Master Beneficiary Record (MBR). Since 1978, the MEF measures uncapped W-2 pre-tax earnings for all Social Security Numbers (SSNs) in the U.S. for each calendar year, with separate information on self-employment and non-self-employment earnings. W-2s are mandatory information returns filed with the Internal Revenue Service for each employee for whom the firm withholds taxes and/or to whom remuneration exceeds a modest threshold. Thus, we have data on earnings regardless of whether an employee files taxes. The data longitudinally follow individuals over time.

The MBR contains information on exact date of birth, exact date of death, month and calendar year of claiming OASI, race, and sex. In the calendar year after an individual dies, earnings and employment appear in the dataset as zeroes; thus, some of an effect on employment could in principle be mediated through an effect on mortality, which would affect the

interpretation—but not the validity—of our results. The effects on earnings that we estimate are nonetheless policy-relevant, in the sense that they reflect the overall effect on employment. We have data on individuals born in the 1918 to 1923 cohorts when they reach ages 50 to 64, in calendar years 1968 to 1987. We classify age in a given calendar year as the highest age an individual attains in that calendar year. Individuals in our cohorts reach ages 63 to 64, the ages at which we investigate the effect of the AET, in 1981 to 1987.

We choose 64 as the oldest age at which to examine employment effects because age 60 and 61 earnings are a better proxy for desired earnings at ages 63 to 64 than for older ages. We cannot use age 62 or 63 earnings as a proxy for desired earnings at age 65 because of potential intensive margin responses to the AET at ages 62 or 63. Moreover, at age 65 individuals with earnings near the under-NRA exempt amount are only exposed to this exempt amount—as opposed to the much higher exempt amount applying to those at NRA and above shown in Figure 1—for only part of the year. This consideration applies *a fortiori* to those over 65.

Several features of the data merit discussion. First, these administrative data are subject to little measurement error. Second, earnings as measured in the dataset are not subject to manipulation through tax deductions, credits, or exemptions, and they are subject to third-party reporting (among the non-self-employed). Third, like most other administrative datasets, the data do not contain information on hours worked, hourly wage rates, amenities at individuals' jobs, underground earnings, assets, savings, or consumption. They also do not contain data on unearned income or marital status.

Table 1 shows summary statistics for our analysis sample. As treatment status depends on earnings, our analysis sample is limited to people with positive base age earnings. The mean yearly employment rate among 50 to 64 year-olds, *i.e.* the percent of the corresponding calendar

years when the individual has positive earnings, is 73.03 percent. Mean earnings (including zeroes) at these ages is $25,617.26 in our main sample. 42.99 percent of the sample is female. For comparison we also show the full sample, not restricted to those with positive base age earnings. We use data on 9,292,092 individuals, corresponding to 87,086,806 observations throughout our sample period. Throughout the paper, all dollar figures are expressed in real 2010 dollars.

V. Results

As a preliminary exercise, in Figure 3 we show the density of earnings at ages 60 and 62, relative to the exempt amount. Figure 3 shows that the density of earnings at age 60 appears smooth near the exempt amount, and that the amount of bunching, calculated using the method of Chetty, Friedman, Olsen, and Pistaferri (2011), is statistically insignificant. This supports the validity of our identification strategy: if the density hypothetically showed evidence of a reaction at age 60 to the future imposition of the AET, this could confound the validity of comparing those under and over the exempt amount at age 60. For comparison, Figure 3 also shows the earnings distribution at age 62, when claimants are subject to the AET. At age 62 we see a markedly different pattern than at age 60, with a large, statistically significant spike in the age 62 earnings density near the exempt amount (as documented in Friedberg 2000 or Gelber, Jones, and Sacks 2013).

Figure 4 shows a graphical depiction of the results. On the x-axis is an individual's age in a base year t. On the y-axis is the probability that they have positive earnings in year t+3. For those with earnings above z^* in year t, the probability of positive earnings falls sharply from t=59 to t=60, exactly the age threshold we would expect if individuals respond to the AET by earning zero once they begin to claim OASI and are subject to the AET. By contrast, for those initially

earning below z^* in year t, the probability of having positive earnings in year t+3 falls much more smoothly, both in percentage point terms (shown in the figure) and in percent terms. This is consistent with the hypothesis that the AET reduces employment, as it has particular "bite" among those with relatively high earnings who are disproportionately subject to the AET. Workers may go back and forth between entering and exiting employment at older ages; our results focus on older workers' employment rates, which represent a key policy concern. One caveat is that our measure of employment may not count a shift from the formal labor market to "off-the-books" employment (Christensen 1990).

Figure 4 shows that the trends in employment for those earning above and below z^* during ages prior to 63 are very similar. Thus, we have reason to believe that anticipatory adjustment to the AET is not a significant issue in our context, as those who are likely to not face the AET have a similar trend in outcomes as those who are most likely to face the AET.

Figure 5 shows the difference in employment rates between the treatment and control groups, at $t + 3$, as a function of age t. After being essentially stable from 50 to 59, the figure shows a sharp decrease from t=59 to t=60; age 60 is exactly when individuals will first be able to show an employment reaction to the AET three years later, when they are age 63. This is followed by another sharp decrease from 63 to 64, consistent with a lagged adjustment to the AET (Gelber, Jones, and Sacks 2013).

Having established these graphical patterns consistent with strong effects of the AET on employment, we now proceed to the regression results. Table 2 shows the regression results corresponding to regression (1) above. In Columns 1 to 3 we estimate a large (in absolute value) and highly statistically significant interaction of the "post" dummy with a dummy for being in the treatment group (i.e. having prior earnings above the exempt amount). This indicates a large

negative effect of the AET on employment of nearly eight percentage points. The estimates grow somewhat from the specification with t=3 to the t=4 and t=5 specifications.

To address the possibility of trends in age suggested by Figure 4, we control for a quadratic trend in age t. Columns 4 to 6 show that in these specifications, we estimate nearly identical results to the corresponding specifications in Columns 1 to 3. This supports the validity of our empirical design.

The next tables show a variety of robustness checks. Table 3 shows that when we use the age range t=55 to t=61, rather than the range t=50 to t=61 in our baseline, we estimate extremely similar results to the baseline. Table 4 shows that when we control for age fixed effects – thus controlling for age in the most flexible possible way – we again estimate nearly identical results to the baseline.

One potential issue is that mean OASI benefits vary between our treatment and control groups. To address this issue, in Table 5 we try controlling for OASI benefits received, interacted with age dummy variables, to allow for differential trends by OASI benefits. Column 1 shows a linear control for OASI benefits; this shows a similar coefficient (-7.155) to the baseline. Columns 2 and 3 show a cubic control for OASI benefits and controls for dummies for each decile of OASI benefits, respectively. These show somewhat smaller, though still large, estimates: the coefficient indicates an effect of 3.947 percentage points in the specification with the cubic control, and 4.996 percentage points in the specification with the decile dummies. Our overall conclusion is that regardless of the specification, the results show large effects of the AET on employment, in the range of several percentage points.

We conduct placebo tests using DD strategies around other "placebo" exempt amounts and find no evidence of effects that are so large and significant. Figure 6 shows two sets of

placebo tests. In Figure 6A, we show that when we implement the same baseline DD strategy as described above, but instead using exempt amounts lower than the true exempt amount (in $25 increments), the placebo estimates we generate are much smaller than the baseline -7.921 percentage point effect at the true exempt amount. Although the placebo point estimates in both of these cases are negative, even subtracting these placebo estimates from our estimate at the true exempt amount would still reveal an employment effect of the AET of several percentage points. This placebo test addresses the concern that there could be differential changes in employment of the magnitude observed in our treatment group relative to the control group, for reasons unrelated to the effect of the AET. For example, this placebo test addresses the concern that differential OASI payments in the treatment and control groups could affect employment through the liquidity channel: the OASI benefit formula replaces average lifetime earnings (*i.e.* Average Indexed Monthly Earnings) with OASI benefits (*i.e.* the Primary Insurance Amount) in a progressive way, implying that OASI benefits rise with current income more steeply at low current income levels than at higher levels. For this reason, we do not run placebos at much higher earnings levels than the true exempt amount – *e.g.* once OASI benefits have been phased out due to the AET – though we do investigate a higher range of earnings in Figure 7.

In Figure 6B, we show that at younger ages not subject to the AET, the placebo estimates we generate around the exempt amount (that applies beginning at age 62) are again much smaller than the baseline -7.921 percentage point effect at the true exempt amount. Specifically, we examine ages t=50 to 58 using the same strategy as in our baseline, but here using ages t=50 to 56 as the "pre" ages, and ages t=57 to 58 as the "post" ages. This set of findings does not address the liquidity concern discussed above, though it does demonstrate that at the exempt amount and

similar placebo exempt amounts, we do not see such striking patterns at younger ages when individuals are not exposed to the AET.

Figure 7 provides a final piece of evidence that we consider to be among our most convincing evidence in this paper that the AET has a strong effect on employment. Figure 7 shows the mean change in the average net-of-tax rate (ANTR) due to the AET for each $2,000 bin of age t distance to exempt amount (Figure 7A), as well as how the estimated treatment effect on employment at age $t + 3$ varies by distance to exempt amount (Figure 7B). The ANTR is defined as the fraction of a dollar that an individual keeps if they are employed rather than non-employed; we calculate this by applying the AET rules to OASI benefits, and calculating the fraction of observed income (*i.e.* earnings plus OASI benefits) that an individual keeps if they choose to earn their desired earnings rather than earn zero. Figure 7A and Figure 7B are strikingly similar. The effect of the AET on the ANTR peaks around $40,000 above the exempt amount because the AET eventually reduces current OASI benefits to zero, after which the effect of the AET on the ANTR begins to fall. The treatment effect increases in magnitude with distance to exempt amount, and is largest in the region in which the AET has the biggest effect on the ANTR. At higher distances above the exempt amount, the treatment effect begins to fall in magnitude, as does the absolute value of the effect on the ANTR. We consider this compelling evidence that the AET has a strong effect on employment; otherwise, it appears very hard to explain why these two graphs track each other (in mirror image) so closely.

The coefficient we estimate has implications for the effect of the AET on the overall employment rate. Our baseline linear specification without controls shows that the AET reduces the employment rate of those in our treatment group by 7.921 percentage points. However, this does not mean that the AET reduces the overall employment rate at ages 63 to 64 by this amount,

for two reasons. First, at ages 60 and 61 averaged, 34 percent of individuals are not employed and therefore are not in our sample. Second, 18 percent of those with positive earnings at age 60 have earnings that are below the exempt amount, implying that they are not subject to the treatment. After deflating due to both of these factors, our estimates imply that the AET reduces employment by 3.7 percentage points in the group we study. If anything, this reflects a lower bound, for example because we observe desired earnings with error.

VI. Conclusion

We show that the AET plays a substantial role in determining older workers' labor force participation decisions, consistent with the results in Gelber, Jones, Sacks, and Song (2017). This indicates that the local estimates in Gelber, Jones, Sacks, and Song (2017) appear to apply to a broader group, in the sense that in a broader sample we still estimate large and highly significant effects of the AET. Indeed, the estimates suggest that the AET reduces the employment rate of the workers we study by 3.7 percentage points. Since the AET applies to those 62 to 66 today, the age range we investigate is relevant to evaluating the effects of the AET as it is currently configured and current policy proposals for changing its parameters. Our results also imply that the planned increases in the Normal Retirement Age, to 67 in 2026, may reduce labor force participation by exposing seniors to the Earnings Test for longer. Our results are consistent with the view that the retirement decisions of the elderly are rather sensitive to incentives (e.g. French 2005).

However, it is important to caution that our results do not necessarily imply that the AET policy is undesirable. By lowering AET benefits earlier in an individual's period of claiming, and thus raising later benefits through benefit adjustment, the AET on average shifts the profile of OASI benefits later in the lifecycle. All else equal, this could lead to welfare gains and more

redistribution in terms of yearly income, as older claimants on average have lower income and assets (as well as worse health) than younger claimants. Evaluating the full welfare and redistributive consequences of the AET, including both the employment and earnings impacts, as well as the impacts through the lifecycle profile of OASI benefits, is an important topic left to future work.

References

Brown, Jeffrey, Arie Kapteyn, Olivia Mitchell, and Teryn Mattox. 2013. "Framing the Social Security Earnings Test." Wharton Pension Research Council Working Paper.

Burtless, Gary, and Robert A. Moffitt. 1985. "The Joint Choice of Retirement Age and Postretirement Hours of Work." *Journal of Labor Economics.* 3: 209-236.

Christensen, Kathleen. 1990. "Bridges Over Troubled Water: How Older Workers View the Labor Market." In *Bridges to Retirement: Trends in the Labor Market for Older Workers* (Peter B. Doeringer, ed.), 175-207. Ithaca, NY: ILR Press, Cornell University.

Diamond, Peter, and Jonathan Gruber. 1999. "Social Security and Retirement in the United States." In *Social Security and Retirement Around the World* (Jonathan Gruber and David A. Wise, eds), 437-473. Chicago: University of Chicago Press.

French, Eric. 2005. "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour." *The Review of Economic Studies*. 72: 395-427.

Friedberg, Leora. 1998. "The Social Security earnings test and labory supply of older men." In *Tax Policy and the Economy* (James M. Poterba, ed.), 121-150. Chicago: University of Chicago Press.

Friedberg, Leora. 2000. "The Labor Supply Effects of the Social Security Earnings Test." *The Review of Economics and Statistics.* 82: 48-63.

Friedberg, Leora, and Anthony Webb. 2009. "New Evidence on the Labor Supply Effects of the Social Security Earnings Test." In *Tax Policy and the Economy* (Jeffrey R. Brown and James M. Poterba, eds.), 1-35. Chicago: University of Chicago Press.

Gelber, Alexander, Damon Jones, and Daniel Sacks. 2013. "Earnings Adjustment Frictions: Evidence from the Social Security Earnings Test." NBER Working Paper No. 19491.

Gelber, Alexander, Damon Jones, Daniel Sacks, and Jae Song. 2017. "Using Kinked Budget Sets to Estimate Extensive Margin Responses: Evidence from the Social Security Earnings Test." NBER Working Paper 23362.

Gruber, Jonathan, and Peter Orzag. 2003. "Does the Social Security Earnings Test Affect Labor Supply and Benefits Receipt?" *National Tax Journal.* 56: 755-773.

Haider, Steven, and David Loughran. 2008. "The Effect of the Social Security Earnings Test on Male Labor Supply: New Evidence from Survey and Administrative Data." *Journal of Human Resources.* 43(1): 57-87.

Liebman, Jeffrey B., and Erzo F.P. Luttmer. 2015. "Would People Behave Differently If They Better Understood Social Security? Evidence From a Field Experiment." *American Economic Journal: Economic Policy* 7(1): 275-299.

Song, Jae G., and Joyce Manchester. 2007. "New evidence on earnings and benefit claims following changes in the retirement earnings test in 2000." *Journal of Public Economics.* 91: 669-700.

Figure 1. *Earnings Test Real Exempt Amount, 1978 to 1987*

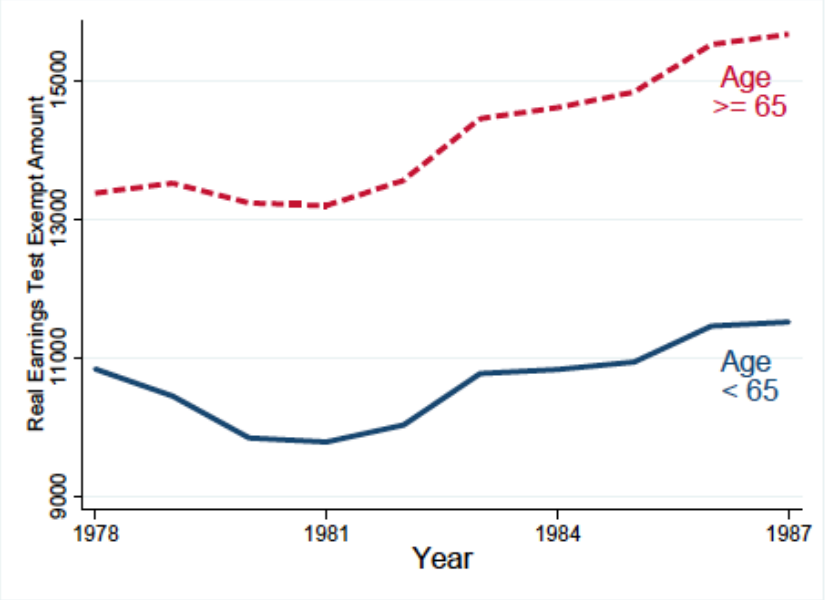

Notes: The figure shows the real value of the exempt amount over time among those 62-64 years old (labeled "Age<65" in the graph) and those 65 and above. The AET applied to earnings of claimants from ages 62 to 71 from 1978 to 1982, but only to claimants aged 62 to 69 from 1983 to 1989. All dollar figures are expressed in real 2010 dollars.

Figure 2. *Histogram of Percent Real Earnings Growth, Ages 59 to 60*

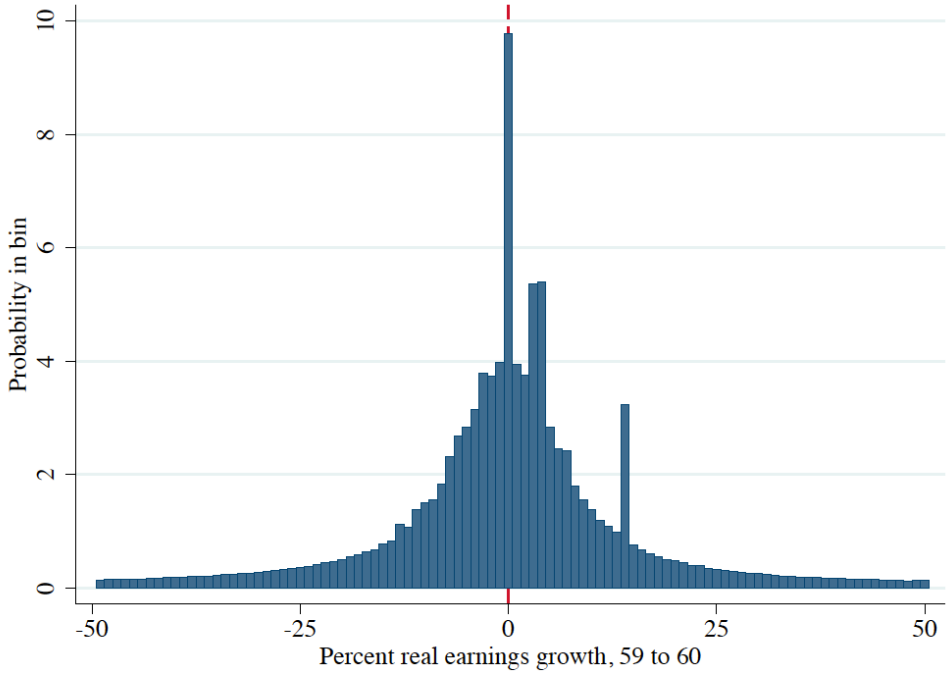

Notes: This histogram shows that there is a spike at zero percent real earnings growth in a "placebo" set of ages, 59 and 60, when individuals do not face the AET. This indicates that a substantial mass of individuals have no growth in desired real earnings, consistent with the assumptions necessary for our RKD to estimate a lower bound on the elasticity as described in the main text. Real earnings in each year are calculated using the CPI-U.

Figure 3. *Earnings Distributions at Ages 60 and 62*

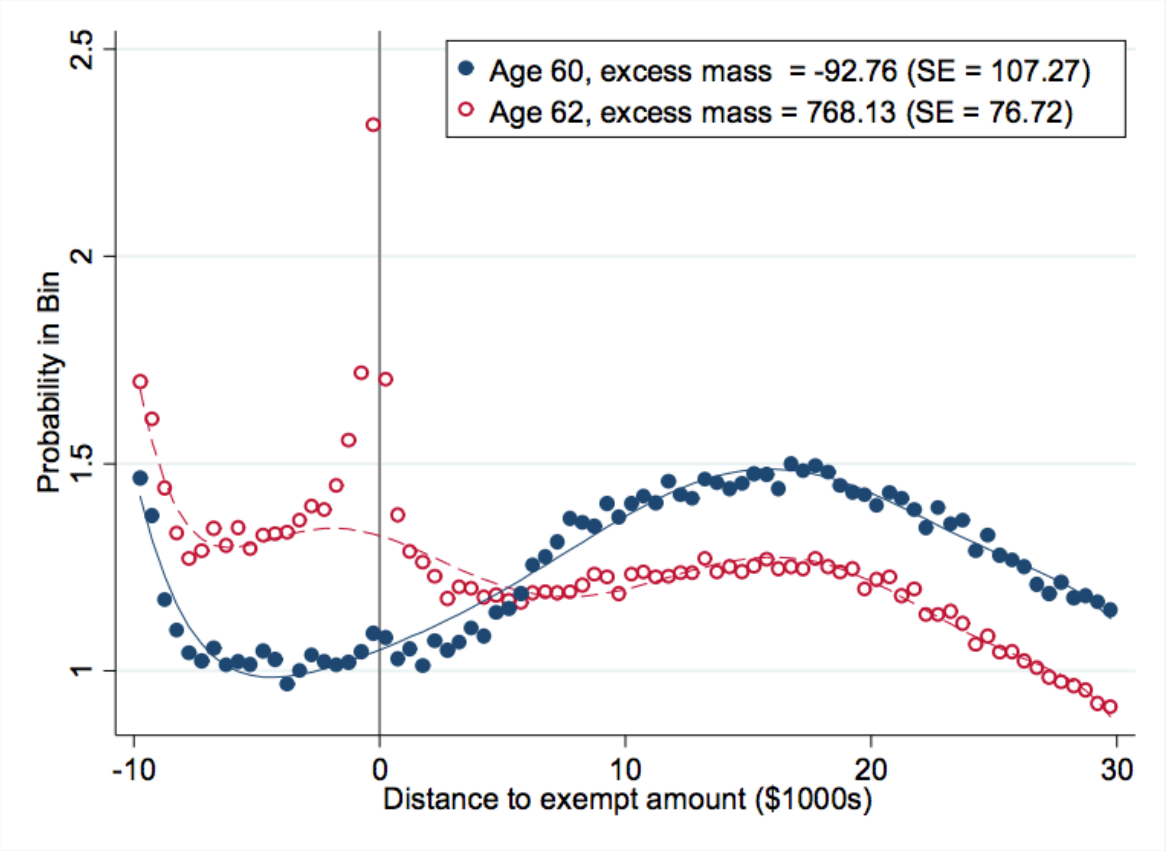

Note: the figure shows the earnings distributions at ages 60 and 62, based on a 25% sample of the SSA data.

Figure 4. *Probability of positive earnings in year t+3, by year-t age and earnings relative to exempt amount*

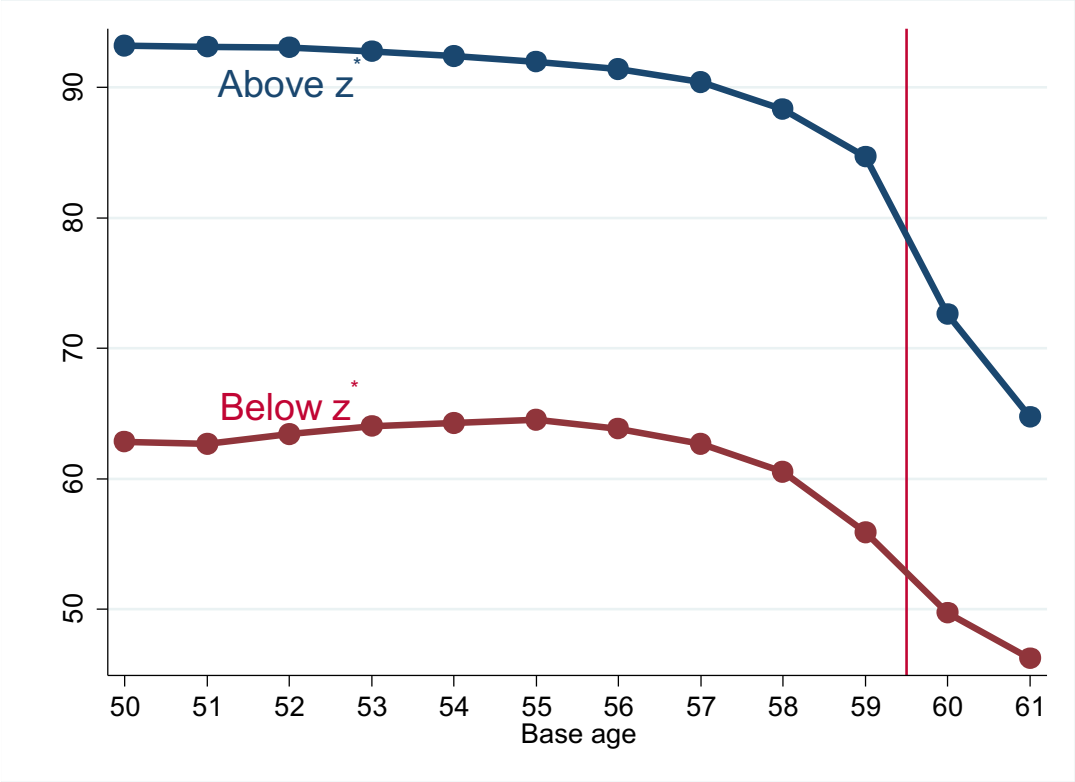

Figure 5. *Difference between probability of positive earnings in year t+3 among those earning above and below the exempt amount, by year-t age*

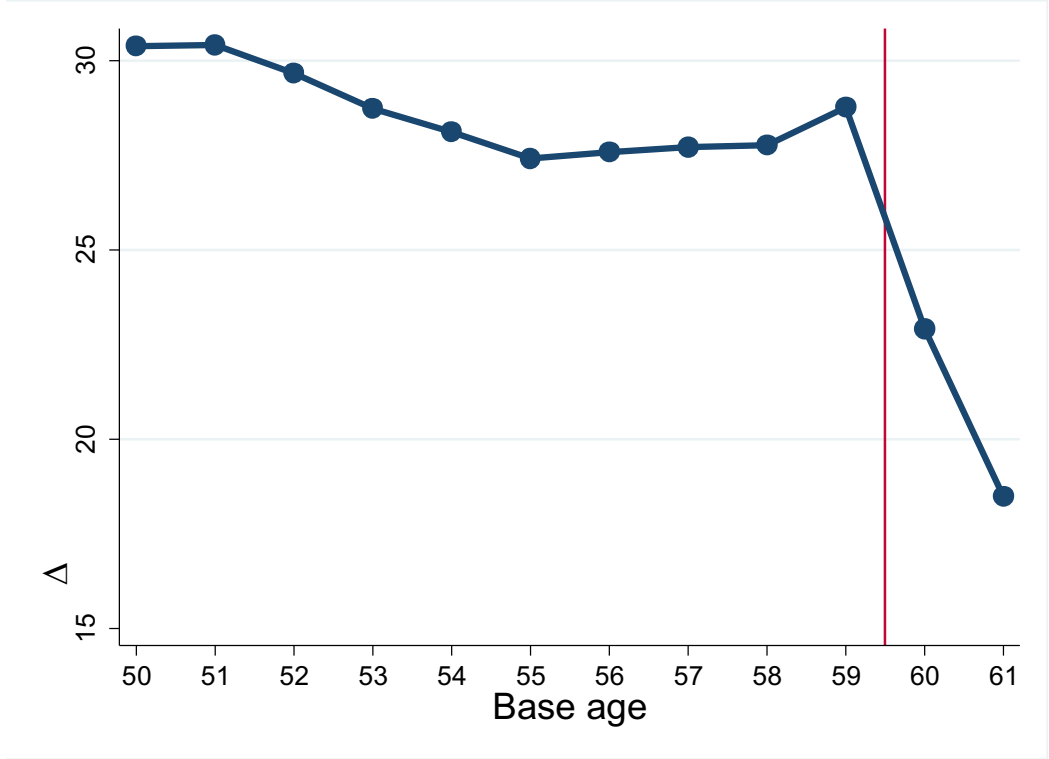

Figure 6. *Placebo estimates*

A. Placebo estimates using lower, placebo "exempt amounts"

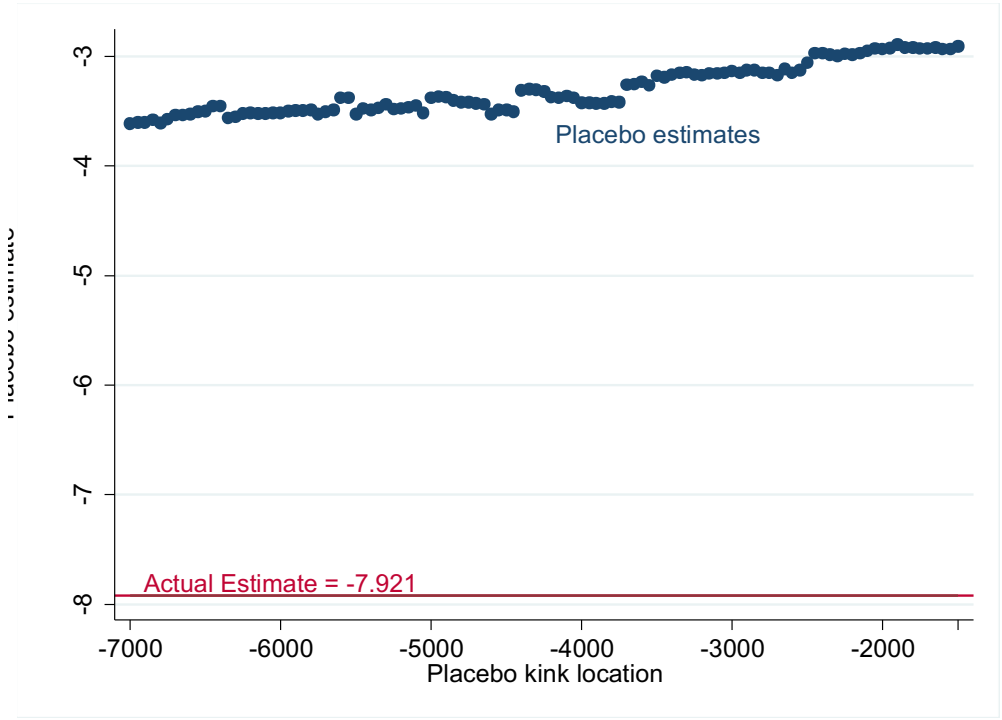

B. Placebo estimates at younger ages

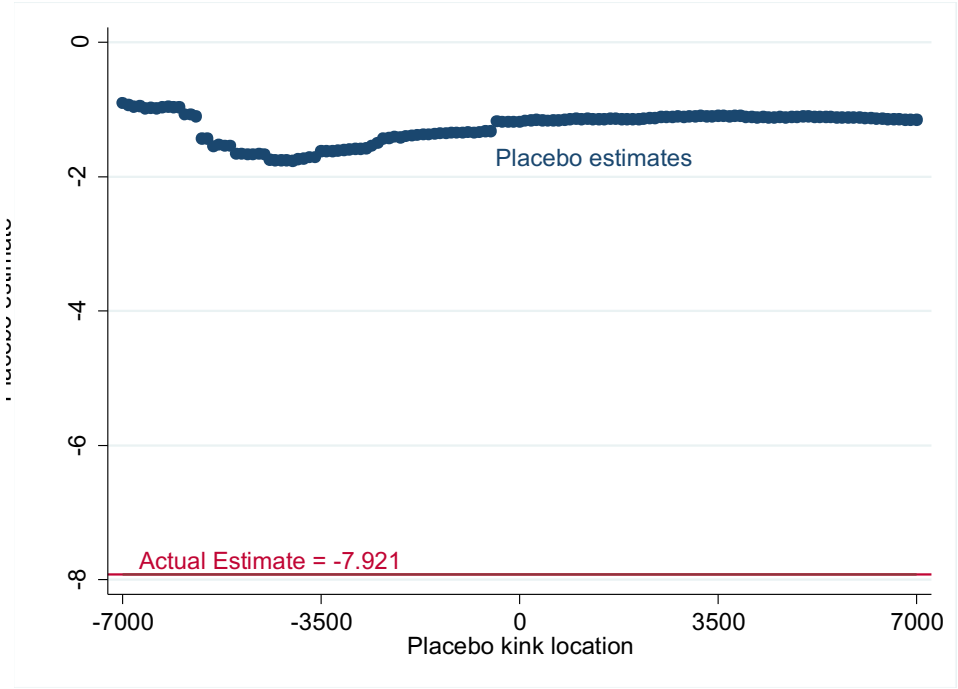

Figure 7. *Change in Average Net-of-Tax Rate Due to the AET, and Treatment Effect, by Base Age Distance to the Exempt Amount*

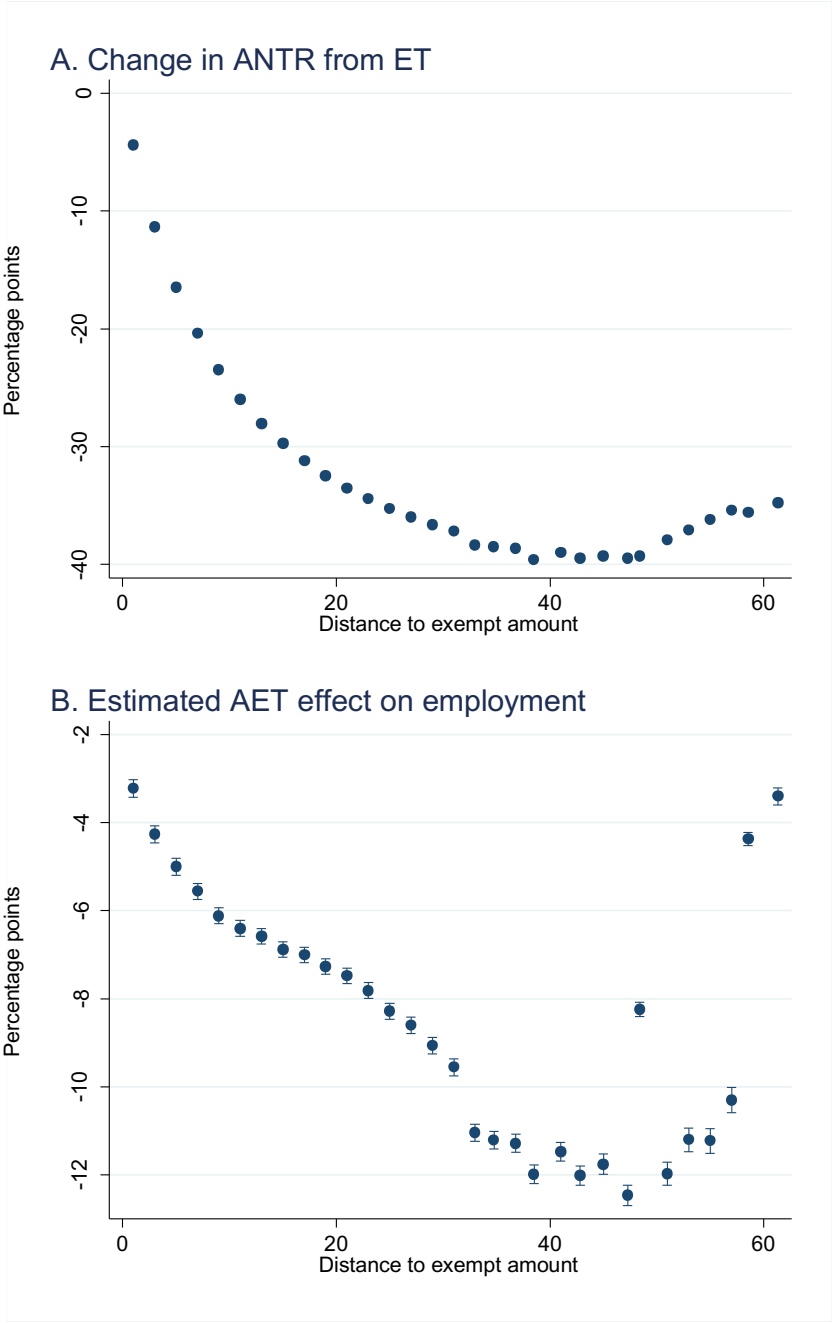

Notes: the figure shows the mean change in the average net-of-tax rate (ANTR) due to the AET for each bin in base year earnings relative to the exempt amount (Figure 6A), as well as how the treatment effect varies by distance to exempt amount (Figure 6B). The ANTR is defined as the fraction of a dollar that an individual keeps if they are employed rather than non-employed. These figures track each other closely. The treatment effect increases with distance to exempt amount, and reaches a peak in the region in which the AET has the biggest effect on the ANTR. At higher distances above the exempt amount, the treatment effect begins to fall, as does the absolute value of the effect on the ANTR.

Table 1. *Summary Statistics: means (standard deviations) of main variables*

	Full Sample	Analysis Sample
Positive earnings dummy, ages 50-64	58.22 (39.69)	73.03 (30.06)
Earnings, ages 50-64	20,396.08 (21,987.37)	25,617.26 (21,769.64)
Claim Age	63.53 (2.26)	63.53 (2.17)
Female dummy	45.55 (49.80)	42.99 (49.51)
White dummy	88.11 (32.37)	89.50 (30.66)
Year of birth	1920.55 (1.70)	1920.53 (1.70)
Number of people	11,676,081	9,292,092

Note: The dummies have been multiplied by 100, so the means reflect percentages.

Table 2. *Main Estimates*

	(1) t=3	(2) t=4	(3) t=5	(4) t=3	(5) t=4	(6) t=5
	Main			With controls		
Post x Treat	-7.921 (0.046)	-10.536 (0.043)	-13.341 (0.042)	-7.872 (0.046)	-10.692 (0.043)	-13.791 (0.042)
# Obs	87,086,806	87,086,806	87,086,806	87,086,806	87,086,806	87,086,806
# People	9,292,092	9,292,092	9,292,092	9,292,092	9,292,092	9,292,092

Notes: The table shows coefficients on "Post x Treat", with standard errors in parentheses. Each of the specifications controls for "Treat" as well as "Post." "Treat" is a dummy if base age earnings are above z^*. "Post" is a dummy if base age is within t years of age 63. The outcome is a dummy for zero earnings (x100) at age a + t, with t ranging from t=50 to t=61. The controls are a quadratic for age t.

Table 3. *Estimates using ages 55-61*

	(1) t=3	(2) t=4	(3) t=5	(4) t=3	(5) t=4	(6) t=5
	Main			With controls		
Post x Treat	-7.066 (0.048)	-9.437 (0.047)	-11.974 (0.047)	-7.157 (0.048)	-9.759 (0.047)	-12.555 (0.047)
# Obs	48,580,452	48,580,452	48,580,452	48,580,452	48,580,452	48,580,452
# People	8,296,628	8,296,628	8,296,628	8,296,628	8,296,628	8,296,628

Notes: The table shows coefficients on "Post x Treat", with standard errors in parentheses. Each of the specifications controls for "Treat" as well as "Post." "Treat" is a dummy if base age earnings are above z^*. "Post" is a dummy if base age is within t years of age 63. The outcome is a dummy for zero earnings (x100) at age a + t, with t as indicated. The controls are a quadratic for age t. Table 3 differs from Table 2 because Table 2 uses the age range t=50 to t=61, whereas Table 3 uses the age range t=55 to t=61.

Table 4. *Controlling for age fixed effects*

	(1) t=3	(2) t=4	(3) t=5	(4) t=3	(5) t=4	(6) t=5
	t=50 to t=61			t=55 to t=61		
Post x	-7.988	-10.785	-13.840	-7.176	-9.751	-12.551
Treat	(0.046)	(0.044)	(0.042)	(0.048)	(0.047)	(0.047)
# Obs	87,086,806	87,086,806	87,086,806	48,580,452	48,580,452	48,580,452
# People	9,292,092	9,292,092	9,292,092	8,296,628	8,296,628	8,296,628

Notes: The table shows coefficients on "Post x Treat", with standard errors in parentheses. Each of the specifications controls for "Treat" as well as dummies for each age. "Treat" is a dummy if base age earnings are above z^*. The outcome is a dummy for zero earnings (x100) at age a + t, with t as indicated.

Table 5. *Controlling for age interacted with OASI benefits*

	(1) Linear	(2) Cubic	(3) Decile dummies
Post x	-7.155	-3.947	-4.996
Treat	(0.050)	(0.052)	(0.051)
# Obs	48,580,452	48,580,452	48,580,452
# People	8,296,628	8,296,628	8,296,628

Notes: The table shows coefficients on "Post x Treat", with standard errors in parentheses. Each of the specifications controls for "Treat" as well as dummies for each age. "Treat" is a dummy if base age earnings are above z^*. The outcome is a dummy for zero earnings (x100) at age a + t, with t as indicated. The specifications differ based on how we control for OASI benefits. In Column 1 we include a linear control for benefits, interacted with each age dummy; in Column 2 we include a set of cubics in OASI benefits, each interacted with each age dummy; and in Column 3 we include dummies for each decile of OASI benefits, each interacted with each age dummy.