The Employment Effects of the Social Security Earnings Test*

Alexander Gelber
UC San Diego 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

November 2018

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 find that the employment rate of those predicted to become subject to the AET decreases substantially relative to those not predicted to become subject to it. The point estimate suggests that the AET reduces the employment rate of Americans aged 63-64 by 3.3 percentage points.

* This research was supported by the U.S. Social Security Administration (SSA) through grant #RRC08098400-09 to the National Bureau of Economic Research as part of the SSA Retirement Research Consortium, as well as by the Alfred P. Sloan Foundation. This research was completed partly while Gelber was on leave at the Stanford Institute for Economic Policy Research, partly funded by a grant from the Alfred P. Sloan foundation. We thank Gary Engelhardt as well as seminar participants at SSA and the Retirement Research Consortium for comments. 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

As the economy-wide employment-to-population ratio has fallen by around four percentage points since its peak in 2000, many policy makers are interested in ways to affect employment rates (Abraham and Kearney, 2017). For older Americans, the Social Security Annual Earnings Test (AET) has a large effect on Social Security Old Age and Survivor Insurance (OASI) benefits and could therefore affect employment substantially. 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 2018, current OASI benefits are reduced by 50 cents for every extra dollar earned above $17,040. The bite of the AET can be substantial. For example, we estimate that in 2003—the last year for which we have the relevant data—the AET reduced the current OASI benefits of 538,000 individuals. The AET reduced their benefits by over half on average (51.4 percent), or $4.3 billion in aggregate.[1]

This large benefit reduction rate may reduce OASI claimants' incentives for additional work. Indeed, the AET is a leading candidate in helping to explain the upward spike in the hazard of retirement at age 62 (Gruber and Wise 1999). The importance of the AET is now increasing as the Normal Retirement Age (NRA) gradually rises from 65 for those born in 1937 and earlier, to 67 for those born in 1960 and later, exposing more OASI claimants to the AET, which does not apply to older ages. Reductions in current benefits due to the AET sometimes

[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 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.

lead to increases in later benefits; nonetheless, as we discuss, several factors may explain why individuals' earnings still respond to the AET.

The prior literature on the AET has tended to focus on the policy's intensive margin effect, *i.e.* the choice of how much to earn, given that an individual earns a positive amount (*e.g.* Burtless and Moffitt, 1985; Friedberg, 1998; Friedberg, 2000; Song and Manchester, 2007; Gelber *et al.*, 2013; Engelhardt and Kumar 2009; Engelhardt and Kumar 2014). 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. This literature generally finds moderate substitution elasticities at the intensive margin with respect to the AET. However, nearly all prior studies find little evidence of extensive margin responses, *i.e.* little evidence that older workers respond to the AET on the margin of whether to work or not. The earlier empirical literature on the AET largely concludes that the policy has little meaningful effect on the labor supply of older men (Viscusi 1979; Burtless and Moffit 1985; Gustman and Steinmeier 1985; Vroman 1985; Honig and Reimers 1989; Leonesio 1990). More recently, researchers have examined the effect of the AET on employment decisions using a difference-in-differences framework. Many studies find little evidence for an effect on the employment rate (Gruber and Orszag, 2003; Song, 2004; Song and Manchester, 2007; Haider and Loughran, 2008), while Friedberg and Webb (2009) find significant but modest effects in some specifications in Current Population Survey data.[2]

In this paper, we revisit the effects of the AET on the employment rate of older workers

[2] Examining the effects of earnings tests in other countries using difference-in-differences designs, Baker and Benjamin (1999) find a significant effect of the Canadian earnings test on weeks worked per year but no significant effect on employment at some point during the year, and Disney and Smith (2002) find inconclusive evidence on the impact of the U.K. earnings test on the employment rate. French (2005) uses method of simulated moments in a lifecycle model to estimate the effects of health, wealth (including all forms of pension wealth, not just incentives created by the AET), and wages on labor supply and retirement; simulations based on these estimates imply that eliminating the AET would cause individuals to retire later on average.

using administrative data and a different methodological approach. 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. At this exempt amount the incentives to work change notably due to the AET. The imposition of the AET above the exempt amount could reduce the incentive to work and therefore reduce employment among those affected in this population. Using a differences-in-differences design, we compare employment rates subsequent to reaching the Social Security retirement age of those with predicted earnings above and below the AET exempt amount, who form the "treatment" and "control" groups, respectively. We implement this design using Social Security Administration (SSA) data on a 25 percent random sample of the U.S. population in birth cohorts 1918 to 1923 over the years 1968 to 1987. Our sample includes over nine million people with over 87 million observations of annual earnings.

Our results show much larger effects on employment than the bulk of 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. This effect occurs suddenly at the ages individuals become subject to the AET, bolstering the credibility of the estimates. Our point estimates suggest that the AET reduces the employment rate of older Americans aged 63-64 by 3.3 percentage points. This conclusion is robust to a battery of placebo and robustness checks.

These findings indicate that the AET is an important factor affecting the work decisions of older Americans and therefore should be a key focus of policy-makers. The combination of a large administrative dataset with individual-level microdata (also used in various forms in Song, 2004, Song and Manchester, 2007, and Haider and Loughran, 2008) and our novel identification strategy leads to estimates of sizeable elasticities. However, the effect we estimate applies to a

younger group than those studied in the difference-in-differences literature cited above: we focus on those locating near the exempt amount in the policy-relevant 63-to-64 year-old group to which the AET currently applies. Our findings are therefore not directly comparable to this previous literature, but provide relatively novel estimates for a younger group of current policy relevance.

Recent work by Gelber *et al.* (2018), 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.[3] Whereas the Gelber *et al.* (2018) study relies on the Regression Kink Design smoothness assumptions, the current paper relies on the differences-in-differences parallel trends assumptions. Gelber *et al.* (2018) study a group that is local to the exempt amount, whereas the current paper's estimates apply both those close to and further from the exempt amount—a broader group of significant policy relevance to evaluating the more general employment effect of the AET.

The paper proceeds as follows. Section II describes the policy environment. Section III describes a framework for interpreting our results. Section IV explains the empirical strategy. Section V describes our data. Section VI proceeds to the results. Section VII 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. To

[3] Because the policy environment, corresponding theoretical framework, data, and outcomes are similar to Gelber *et al.* (2018), the corresponding sections of this paper have overlap with that previous paper.

illustrate how the AET works, consider a 63-year-old earning $23,040 in 2018, receiving $1,000 in monthly benefits, and facing a $17,040 exempt amount and a 50 percent benefit reduction rate (BRR). Her current annual benefits would be reduced by $3,000=($23,040-$17,040)×50%, equal to 3 months of benefits. If that same worker instead earned $50,000, annual benefits would be reduced by $12,000, rather than being reduced by $16,480=($50,000-$17,040)×50%, because benefits are not reduced below zero.

Both the exempt amount and the BRR have changed over time. In Figure 1 we show the exempt amount for the years we study. For those whose age is at the NRA and over, the exempt amount is substantially higher than for those below the NRA. Over the main years we study, 1981 to 1987, the exempt amount for those under NRA ranged from $9,787 to $11,517, while the exempt amount for those NRA and over was on average around $3,800 higher than the amount for those under NRA. For those under the NRA but above the EEA—the main group that our empirical work studies—the (BRR) was 50 percent throughout the period we study, 1981 to 1987. During our period, the AET applied to earnings from ages 62 to 71 during 1981 to 1982, and from ages 62 to 69 during 1983 to 1987.

During the period we study, the AET rules for married couples imply that the couple's total benefit is reduced at the rate of one dollar for every two dollars that each spouse's individual current earnings exceed the exempt amount applying to each spouse. How this occurs depends on how benefits are claimed. If each spouse is a primary OASI beneficiary, then the AET reduces each spouse's separate benefit by the BRR multiplied by that spouse's individual current earnings in excess of the exempt amount. If one spouse is a primary beneficiary and the other is a secondary or dual-entitled beneficiary, the couple's total benefits are reduced by the BRR multiplied by the primary beneficiary's current individual earnings in excess of the exempt

amount, and further reduced by the BRR multiplied by the secondary or dual beneficiary's current individual earnings in excess of the exempt amount. In either case, the relevant amount for applying the AET is each individual's current annual earnings, which we observe in our data.[4] Because the couple's total benefit is reduced at a 50 percent BRR as described above, this allows us to interpret our estimates within a unitary model (Becker 1976) as the effect of the AET applying to a given spouse's earnings (but potentially reducing both spouses' combined current OASI benefits at the BRR) on that spouse's employment.

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. Returning to the example above, consider the 63-year-old receiving $1,000 in monthly benefits due to the AET. Upon reaching the NRA, her monthly benefits would increase by around $16.50 = 0.0055 x 3 x $1,000. On average, this adjustment is roughly actuarially fair when considering the timing of claiming OASI (Diamond and Gruber, 1999).

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

[4] Due to the nature of our data (described in detail later), 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.

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. Framework for interpreting the empirical results

In this section, we sketch a basic framework that is helpful in interpreting the empirical strategy and results. We model how the AET impacts an individual's decision of whether or not to have a positive amount of earnings, which we refer to as the "employment" decision. Throughout, we make use of a potential outcomes framework (Rubin, 1974). We index two potential states of the world by $j \in \{0,1\}$.

To capture the real-world features of how OASI benefits, taxes, and the AET work, our framework incorporates all three.[5] Following previous literature (*e.g.* Friedberg, 1998; Friedberg, 2000), we model the AET as creating a positive benefit reduction rate (BRR) for some individuals above the exempt amount, consistent with the empirical finding in this previous literature that some individuals bunch at the exempt amount. Individuals receive a level of current benefits that is potentially a function of earnings, *i.e.* $B_j(z)$, where $B_j(\cdot)$ denotes their current benefit in state *j,* and *z* denotes their pre-tax and pre-benefits earnings. The "pre-reduction" level of benefits is *b*, which refers to the OASI benefits received before accounting for the effects of the AET or taxes.[6] Current benefits, $B_j(z)$, are determined both by *b* as well as

[5] It would alternatively be possible to model the effects of OASI benefits and taxes using a single function, but we have modeled them separately to capture the reality of how the tax system and the AET operate separately. They are administered by separate agencies: the Internal Revenue Service and the Social Security Administration, respectively.

[6] For notational simplicity, we have made the benefit constant across individuals. In reality, each individual may receive a different level of pre-reduction benefits. The main issue this affects for our purposes is the earnings level at which the benefit is phased out entirely, which in reality can be different across individuals.

by any reductions in benefits due to the AET. Finally, there is a linear tax on earnings, *i.e.* $T(z) = \tau_0 z$, which does not vary by state.[7] This tax, which reduces net earnings relative to gross earnings, is separate from the AET, which only acts to reduce OASI benefits.[8] Total post-tax and post-benefit resources are therefore:

$$z - T(z) + B_j(z) = (1 - \tau_0)z + B_j(z).$$

In state 0, when there is no AET, the current benefit level is independent of earnings, *i.e.* $B_0(z) = b$. Therefore, individuals face a flat net "benefit reduction rate." That is, as earnings increase, the marginal reduction in post-tax and post-benefit resources is simply τ_0 . In state 1, the AET BRR is τ_b where the AET reduces benefits at the margin, *i.e.* for earnings above the exempt amount but below the point at which benefits have been phased out entirely. The presence of the AET introduces two changes in slope to the budget set, one at z_1^* and another at $z_2^*(b)$, due to reductions in current benefits:

$$B_1(z) = \begin{cases} b \text{ for } z \leq z_1^* \\ b - \tau_b(z - z_1^*) \text{ for } z_1^* < z \leq z_2^*(b) \\ 0 \text{ for } z_2^*(b) < z \end{cases}$$

The first change in slope occurs at the point at which the AET is imposed, z_1^*, while the second change in slope occurs at the point where OASI benefits are phased out entirely, $z_2^*(b)$. At the higher amount, $z_2^*(b)$, the net benefit reduction rate returns to its lower level, creating a non-convex kink in the budget set. This second threshold is a function of OASI benefits, and is defined as follows:

$$z_2^*(b) = z_1^* + b/\tau_b.$$

[7] We do not model taxes on Social Security benefits for simplicity; adding taxes on benefits would not change the qualitative predictions of the framework. Social Security benefits were untaxed until 1984, so benefits escaped taxation fully in most of our sample years (from 1978 to 1983). Starting in 1984, benefits were only taxed above an income threshold that was well above the AET exempt amount, implying that most benefits still escaped taxation.
[8] Introducing non-linear taxes for each individual would not qualitatively affect the predictions of this section as long as they are linear on average in the relevant range.

This second threshold varies at the individual level, based on the size of one's OASI annual benefit.

Following previous literature, we assume individuals have a smooth distribution of "ability," which governs the tradeoff between leisure and consumption. In the presence of a linear tax, this should result in a smooth distribution of earnings conditional on working (*e.g.* Hausman, 1981; Saez, 2010; Kleven and Waseem, 2013). In the presence of the AET, a standard model predicts an intensive margin response with excess mass in earnings, or "bunching," to be present at the convex kink created at z_1^* (Gelber *et al.* 2013). At the extensive margin, to capture the realistic pattern of potential entry to or exit from non-trivial levels of earnings, we can assume a fixed cost of employment (Cogan, 1981; Eissa *et al.*, 2008). In this case, extensive margin decisions are a function of the average net-of-benefit-reduction rate (ANBRR), defined as $ANBRR \equiv 1 - \frac{[(T(z)-B(z))-(T(0)-B(0))]}{z}$.[9] The ANBRR reflects the fraction of an individual's gross income that s/he keeps, net of both taxes and benefits, if s/he is employed at earnings level *z* rather than earning zero.

To demonstrate the impact of a kink on the decision to work in this context, we illustrate the extensive margin incentives created by the AET in Figure 2. Here we plot the ANBRR as a function of counterfactual earnings, that is, earnings conditional on working, in the counterfactual state where there is only a linear tax. We denote these potential earnings as $\tilde{z}_0$. We distinguish between these earnings and *realized* earnings, *z*, which incorporate the extensive margin decision and can be zero. The ANBRR measures the share of pre-tax income that is kept after taxes when working and earning *z*. In state 0, the ANBRR is constant at $1 - \tau_0$. This is

[9] Gelber *et al.* (2018) give an extended discussion of extensive margin earnings decisions in the presence of a kinked budget set.

represented by a dashed line. The solid line in Figure 2 shows that with the nonlinear budget set created by the AET, the ANBRR is $1-\tau_0$ below z_1^*, but becomes $1-\tau_0-\tau_b(\tilde{z}_0-z_1^*)/\tilde{z}_0$ above z_1^*, and therefore begins to decrease in $\tilde{z}_0$. However, after the benefit has been entirely phased out, the ANBRR becomes $1-\tau_0-b/\tilde{z}_0$ at $z_2^*(b)$, begins to increase in $\tilde{z}_0$, and eventually asymptotes back to $1-\tau_0$ for large enough $\tilde{z}_0$.

Figure 3 illustrates how extensive margin decisions respond to the introduction of the AET, as a function of counterfactual earnings. The *x*-axis measures counterfactual earnings in state 0 conditional on having positive earnings, $\tilde{z}_0$. The *y*-axis plots an illustrative employment rate. The dashed line represents a presumed smooth relationship between the employment rate in state 0 under a linear tax schedule, *i.e.* $\Pr(z_0>0|\tilde{z}_0)$, and earnings conditional on having positive earnings, *i.e.* $\tilde{z}_0$. Following the pattern of the ANBRR, the probability of positive earnings begins to diverge at the kink at z_1^*. This pattern motivates our difference-in-difference empirical strategy: people with desired earnings below the exempt amount are essentially not subject to the earnings test, but people with desired earnings above the exempt amount are. Around the non-convex kink, $z_2^*(b)$, the ANBRR begins to increase again and the difference between the two employment rates begins to decrease.[10]

The discussion here focuses on the broad pattern of employment, driven by the initially increasing and subsequently decreasing difference in ANBRR between the linear tax schedule and the nonlinear tax schedule. This is different from Gelber *et al.* (2018), who focus on the patterns of employment in a small neighborhood around the first kink at z_1^*. Gelber *et al.* (2018)

[10] More precisely, if the marginal effect of the ANBRR on the employment rate does not depend on counterfactual earnings, the magnitude of the effect of the AET on the employment rate should decrease beginning at the non-convex kink. However, if the marginal effect of the ANBRR on the employment rate varies depending on counterfactual earnings, then the effect of the AET on the employment rate may be maximized at a counterfactual earnings level other than the level associated with the non-convex kink. In either case, the effect of the AET on employment follows a U-shaped pattern, as Figure 3 illustrates.

show that the employment rate, $\Pr(z_{n1} > 0|\tilde{z}_0)$, should exhibit a discontinuous change in slope at z_1^* when individuals are unconstrained at the intensive margin in moving their earnings to the kink z_1^*. The results of Gelber *et al.* (2018) helps us to interpret our results in this paper. As they discuss, those observed to bunch are unconstrained, while others appear to be constrained because Gelber *et al.* (2018) do find a discontinuous change in slope at z_1^*. Thus, as in their paper, the elasticity we estimate in this paper should be interpreted as an "observed" elasticity (in the terminology of Chetty 2012) reflecting a weighted average of the responses of those constrained and unconstrained from bunching at the exempt amount (Gelber *et al.* 2018).

B. Specific features of the policy setting

Our policy setting has a number of specific features that can be incorporated into the framework. First, individuals are subject to the AET only when they choose to claim OASI; this mirrors other transfer programs that must be claimed to receive benefits. However, the results should still be interpreted as reflecting an observed elasticity. Our observed elasticity remains of interest regardless of the effect on claiming, in the sense that policy-makers are interested in the raw employment effects of changing the AET.

Second, we considered a static setting in which individuals only consider the current period's budget set and incentives. In a dynamic setting, we may observe individuals serially facing a linear tax (corresponding to state 0), followed temporally by a tax that creates a convex kink (corresponding to state 1); indeed this is the case in our empirical application, in which individuals are not subject to the AET initially and become subject later. In Gelber *et al.* (2018), we present a fully dynamic, multi-period model with a joint decision over saving and earnings; we show that under a set of empirically relevant assumptions our results still hold, and we discuss the interpretation of our estimated results in a dynamic model. If agents act as if they do

not anticipate the change in an initial period—consistent with the empirical findings in Gelber *et al.* (2013, 2018)—then we can interpret the results as reflecting the impact of an unanticipated change in policy. Third, Gelber *et al.* (2018) also extend this basic dynamic model to tailor it to our particular policy setting by showing that our method still applies when reductions in current benefits due to the AET can lead to increases in later benefits, as under benefit enhancement.

IV. Empirical strategy

A. Regression specification

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. We define "potential" earnings as the earnings individuals would earn in the absence of the incentives created by the AET (Gelber, Jones, and Sacks 2013; Gelber *et al.* 2018).

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 or 4. We explain below why we consider a minimum lead of 3. We denote the AET exempt amount by z^* and compare the difference in the probability of positive earnings for those with $z_{ia} \geq z^*$ to those with $z_{ia} < z^*$ when $a + t < 63$ and when $a + t \geq 63$.

Thus, our approach represents 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 $a + t \geq 63$.[11]

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

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

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 consider $t = 3$ or $t = 4$ 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 fixed effect capturing the effect of having base-period (*i.e.* age a) earnings above z^*. 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, and relative to the pre-age 63 difference. We cluster our standard errors at the individual level to account for intra-individual correlations, such as those arising from serial correlation. We also perform two "permutation" tests from placebo designs to generate an alternative set of standard errors, which generate comparable conclusions.

B. Validity and implementation of method

We treat age 63 as the "post period" because that is the first year one would expect to observe an extensive margin effect of the AET. The AET first applies to claimants when they

[11] 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.

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.

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 remain constant across ages. 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 earning a positive amount in the counterfactual scenario where the AET is not in effect. We would ideally model the probability of 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 decrease in response to the AET.

Many other papers have grappled with the issue of how to proxy for earnings or wages if individuals choose to work, and thus how to proxy for the incentive to work. Given that the econometrician does not directly observe counterfactual earnings, there is no alternative to making some assumption. One solution is a selection correction in the context of the effect of wages on labor supply, which generally requires functional form assumptions (Heckman, 1979) or very powerful instruments (Powell, 1994). Another solution is to use demographics to impute earnings if an individual works (*e.g.* Meyer and Rosenbaum, 2001), which is more difficult in settings such as ours with a limited number of demographic variables in our administrative data.

To circumvent this problem of endogenous earnings decisions at the ages directly subject to the AET, we will use lagged earnings as a proxy for potential earnings in the present period. 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, 2018).

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 subsequent period exhibits a spike at zero. Figure 4 shows that from age 59 to age 60—a placebo set of ages during which our sample is not subject to the AET—a large mass of real earnings growth does occur at or near zero percent growth.

While investigating a lag of *t=4* as a robustness check, in the baseline we choose to use a three-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). 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). A further advantage of using a 3-year lag (rather than 2) is that we can examine age 64 employment (which would be impossible with a 2-year lag because age 62 earnings respond to the AET). As our baseline we use the shortest lag possible, *i.e.* $t=3$, that satisfies these conditions. We avoid using a lag that is greater than $t=4$ because this would involve examining employment outcomes at ages above 65, when the exempt amount is much higher.

Although we use age 60 earnings to proxy for desired earnings at ages 63 to 64 as above, our difference-in-difference strategy effectively uses the employment rate at ages 63 to 64 for those with age 60 earnings under the exempt amount to reflect state 0 employment rates, while the employment rate at ages 63 to 64 for those with age 60 earnings over the exempt amount reflects state 1 employment rates. Given that age 60 earnings (among those who are employed at age 60) measure state 0 earnings with error, our empirical strategy should yield a lower bound on employment responses (see Gelber *et al.* 2018 for a detailed discussion in this context).

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

Second, if those with prior 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.

V. 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.

Due to institutional considerations and computational constraints, we focus on a specific set of cohorts observed over certain calendar years. We focus on a sample that is subject to the AET in 1978 and after. The reason is that we observe only calendar year earnings, and the AET has effectively applied to calendar year earnings beginning in 1978 (Gelber *et al.*, 2013, Appendix A). Before 1978, the AET was applied to quarterly earnings, which we do not observe in the data. This limits our focus to individuals born in 1918 or later, who turned age 60 in 1978 or later. Due to computational constraints, we were able to obtain data on individuals born from 1918 until 1923, and due to these constraints it was also necessary to draw a 25 percent random sample of this group. In our baseline we observe these individuals until they turn 64, which occurs as late as 1987 (in the 1923 birth cohort). 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 main ages at which we investigate the effect of the AET, in 1981 to 1987.

We choose 64 as an oldest age at which to examine employment effects because age 60 earnings are a better proxy for desired earnings at ages 63 to 64 than for older ages. We cannot use earnings at ages 62, 63, or older as a proxy for desired earnings at even older ages because of potential intensive margin responses to the AET once individuals have reached EEA. 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. In our broadest sample, 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.

VI. Results

A. Graphical patterns

As a preliminary exercise, in Figure 5 we show the density of earnings at ages 60 and 62, relative to the exempt amount. Figure 5 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 5 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 6 shows an initial graphical depiction of the results. On the *x*-axis is an individual's age in year $a+3$, where a is the base age. We refer to age $a+3$ as the "outcome age." On the *y*-axis is the probability that they have positive earnings at the outcome age. For those with earnings above z^* in year a, the probability of positive earnings falls sharply and substantially from outcome ages 62 to 63, 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 a, the probability of having positive earnings at the outcome age falls to a much smaller extent, both in percentage point terms (shown in the figure) and in percent terms, and much less sharply (relative to the pre-trend) than for those initially earning above z^* in year a.[12] 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.

[12] It is not surprising that employment falls, albeit relatively smoothly, from ages 62 to 63 even for those initially earning below z^*; employment rates gradually fall at older ages (e.g. Maestas 2010). Moreover, pension programs could have income effects on employment that reduce employment substantially (Fetter and Lockwood forthcoming).

Figure 6 shows that the trends in employment for those earning above and below z^* during outcome 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 7 shows the difference in employment rates between the treatment and control groups at the outcome age, as a function of age. After essentially remaining stable from outcome ages 53 to 62, the figure shows a sharp decrease from 62 to 63; age 63 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 a further substantial decrease from 63 to 64, consistent with a lagged adjustment to the AET (Gelber, Jones, and Sacks 2013, 2018).

B. Regression Results

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 our baseline, we investigate the employment outcome at ages 58 to 64, to focus near the sharp change in incentives between ages 62 to 63. 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 effect of the AET on employment of -7.066 percentage points. The estimates grow somewhat from the specification with t=3 to the t=4 specification.[13]

[13] In the $t=4$ specification, we use the same set of base ages, 55-61, as in the baseline $t=3$ specification, to hold this set of ages constant across specifications. Thus, the ages when we observe the outcome reach as high as 65 (where 65=61+4). Recall that at age 65 individuals are exposed to the under-NRA exempt amount—as opposed to the much

To address the possibility of a trend in age that is both non-linear and common to the treatment and control groups, we control for a quadratic trend in age *t*. Columns 3 to 4 show that in these specifications, we estimate very similar results to the corresponding specifications in Columns 1 to 2. This supports the validity of our empirical design. In Appendix Table 1, we show that the results are little changed by additionally controlling for gender, race, and year of birth, and we further show that the results are similar when controlling for linear trends in age that are separate for the treatment and control groups.

The next tables show a variety of further robustness checks. Table 3 shows that when we use the age range *t*=50 to *t*=61 shown in Figures 6 and 7, rather than the range *t*=55 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.[14] Column 1 shows a linear control for OASI benefits; this shows a similar coefficient (-7.448) 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 columns show somewhat smaller, though still large, estimates: the coefficient indicates an effect of -5.705 percentage points in the specification with

higher exempt amount applying to those at NRA and above shown in Figure 1—for only part of the year. Since age 65 is intermediate in this sense, it makes sense to show the results both with this age included (to hold the set of base ages constant across the specifications), and alternatively without this age included (to hold the set of outcome ages constant across the specifications). When excluding age 65, the coefficient estimate is -8.27, which is in the same range as the comparable coefficient estimate in Table 2 Column 2 from a specification that includes age 65.

[14] Although we lack a measure of OASI benefits in each year in the SSA administrative data, we have information in the SSA data on the benefits the individual received in the final year that the individual is recorded receiving OASI in the data. For the purpose of these regressions, we deflate this measure of benefits to the appropriate year by CPI-U to form a proxy for benefits received in this year.

the cubic control, and -6.403 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 this sample, 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 8 shows two sets of placebo tests. In Figure 8A, we show that when we implement the same baseline DD strategy as described above, but instead using "placebo exempt amounts" substantially lower than the true exempt amount (in $25 increments), the placebo estimates we generate are only a small fraction of the baseline estimate 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.[15]

In Figure 8B, 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

[15] For this reason, we do not run placebos at 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 9. We also do not use placebo exempt amounts more than $7,000 below the true exempt amount; we cannot go too far below the exempt amount to avoid falling below a placebo exempt amount of $0.

than the baseline effect at the true exempt amount. Specifically, we examine ages t=50 to t=58 using the same strategy as in our baseline, but here using ages t=50 to t=56 as the "pre" ages, and ages t=57 to t=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.[16]

It is possible to use the variances of the estimates in the placebo tests generate alternative standard errors on our coefficients, in the spirit of a "permutation" test, rather than the standard errors shown in the tables. These show remarkably similar standard errors across the two sets of placebo tests that we run. When we use the variance of the estimates shown in the placebo test in Figure 8A, based on placebo exempt amounts below the true exempt amount, the implied standard error is 0.219. Alternatively, when we use the variance of the estimates shown in the placebo test in Figure 8B, based on placebo exempt amounts at younger ages, the implied standard error is nearly identical, 0.229. Whether using these standard errors or those shown in Table 1, we conclude that our estimates are highly statistically significant.

Figure 9 provides a final piece of evidence that the AET has a strong effect on employment. Figure 9 shows how the estimated treatment effect on employment at age $t+3$ varies by distance to exempt amount. The treatment effect increases in absolute magnitude with distance to exempt amount. At higher distances above the exempt amount, the treatment effect begins to fall in absolute magnitude.

[16] In this case we are not concerned about the liquidity effects of OASI, as younger ages do not actually receive OASI, so we examine a range of placebo exempt amounts from $7,000 below to $7,000 above the true exempt amount.

At a broad level, this is the U-shaped pattern we would expect to see, given the U-shaped incentives that the AET creates. The AET reduces current benefits at the margin, with a 50 percent BRR, until current benefits are phased out entirely and reach zero. As a result, the ANBRR should slowly fall above the exempt amount, as the absolute effect of the AET on the ANBRR increases. However, after current earnings have reached the level where current benefits are phased out entirely, $z_2^*(b)$ in the notation of our model, the effect of the AET on the ANBRR should fall, as the AET's effect on benefits becomes smaller relative to current earnings.[17] We view this U-shaped pattern as helpful confirmation that the employment effects we estimate are due to the effects of the AET, which could be expected to generate exactly such a pattern of effects. It is arguably difficult to construct alternative explanations for this U-shaped pattern.

Our finding of a substantial employment response to the AET complement those in Gelber *et al.* (2018), who focus on patterns within $3,000 of the exempt amount, and therefore leave open the questions of whether the results generalize to a broader group and what the overall effect on employment is in this age range. Thus, the current study examines a broader group of interest to policy-makers, and indicates that the local estimates in Gelber *et al.* (2018) 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. In the online Appendix, we briefly recapitulate some of the key findings from Gelber *et al.* (2018), which bolster the credibility of the current results further.

D. Implications for the employment effects of the AET

[17] To determine whether the absolute effect of the AET on employment is maximized around the same counterfactual earnings level that the absolute effect of the AET on the ANBRR is maximized, we would need data on spousal benefits; as described above, when one spouse is a dual or secondary beneficiary, the AET reduces the combined benefits of both spouses at the BRR. Our data do not contain information on spousal benefits.

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 7.066 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, 45 percent of individuals are not employed and therefore are not in our sample. Second, among employed 60-61 year-olds in our sample, 16 percent of have earnings that are below the exempt amount, implying that they are not subject to treatment. After deflating due to both of these factors, our estimates imply that the AET reduces employment by 3.3 percentage points in the group we study. If anything, this reflects a lower bound, for example because we observe desired earnings with error.[18]

VII. Conclusion

We show that the AET plays a substantial role in determining older workers' labor force participation decisions, reducing the employment rate of the workers we study by 3.3 percentage points. Our results are more generally consistent with the view that the retirement decisions of older individuals are rather sensitive to incentives (*e.g.* Laitner and Silverman 2012).

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. Moreover, in the 63-to-64 year-old age group the AET operates in a comparable today as it did in the 1980s; in this age range the AET still imposes a 50 percent benefit reduction rate on current earnings and has the same actuarial adjustment of subsequent benefits. While our paper studies employment outcomes in the late 1970s and 1980s, in the 1980s older American's employment rates were several percentage points lower than today

[18] 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).

(Goldin and Katz 2018); with a similar elasticity one might expect still larger percentage point effects on employment in the current environment with higher baseline levels of employment. 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. If individuals react to Social Security incentives only after learning about them (Gelber, Jones, and Sacks 2013; Armour and Lovenheim 2016), then the impacts we find could if anything increase over time at older ages as more individuals learn about the Earnings Test over time after ages 63-64.

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 if the optimal time profile of Social Security benefits slopes upward (Feldstein 1990), 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—and lower poverty rates among older Americans (see Figinski and Neumark 2016 on older women). Nonetheless, the substantial employment effects we have documented in our work represent an important input into a calculation of the welfare effects of the AET, in particular by suggesting important distortions from the AET. 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

Abraham, Katherine, and Melissa Kearney. 2017. "Explaining the Decline in the U.S. Employment-to-Population Ratio: A Review of the Evidence." NBER Working Paper 24333.

Armour, Philip and Michael F. Lovenheim. 2016. "The Effect of Social Security Information on the Labor Supply and Savings of Older Americans." University of Michigan Retirement Research Center (MRRC) Working Paper, WP 2016-361.

Baker, Michael, and Daniel Benjamin. 1999. "How do retirement tests affect the labor supply of older men?" *Journal of Public Economics*, 71, 27-51.

Becker, Gary. 1976. "Altruism, Egoism, and Genetic Fitness: Economics and Sociobiology." *Journal of Economic Literature*, 14(3), 817-826.

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.

Chetty, Raj. 2012. "Bounds on Elasticities with Optimization Frictions." *Econometrica*, 80(3), 969-1018.

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.

Cogan, John. 1981. "Fixed Costs and Labor Supply." *Econometrica*, 49, 945-963.

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.

Disney, Richard, and Sarah Smith. 2002. "The labour supply effect of the abolition of the earnings rule for older workers in the United Kingdom." *Economic Journal*, 112, C136-C152.

Eissa, Nada, Henrik Kleven, and Claus Kreiner. 2008. "Evaluation of four tax reforms in the United States: Labor supply and welfare effects for single mothers." *Journal of Public Economics*, 92(3-4), 795-816.

Engelhardt, Gary, and Anil Kumar. 2009. "The Repeal of the Retirement Earnings Test and the Labor Supply of Older Men." *Journal of Pension Economics and Finance* 8(4), 429-450.

Engelhardt, Gary, and Anil Kumar. 2014. "Taxes and the Labor Supply of Older Americans: Recent Evidence from the Social Security Earnings Test." *National Tax Journal* 67(2): 443-458.

Feldstein, Martin. 1990. "Imperfect Annuity Markets, Unintended Bequests, and the Optimal Age Structure of Social Security Benefits." *Journal of Public Economics* 41(4): 31-43.

Fetter, Daniel, and Lee Lockwood. "Government Old-Age Support and Labor Supply: Evidence from the Old Age Assistance Program." Forthcoming, *American Economic Review.*

Figinski, Theodore, David Neumark, 2018. "Does Eliminating the Earnings Test Increase the Incidence of Low Income Among Older Women?" *Research on Aging*, 40(1): 27-53.

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. 2018. “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.

Gruber, Jonathan, and David Wise. 1999. "Introduction and Summary." In Gruber, Jonathan, and David Wise, *eds*., Social Security and Retirement Around the World (pp. 1-35). Chicago: University of Chicago Press.

Gustman, Alan, and Steinmeier, Thomas. 1985. “The 1983 Social Security Reforms and Labor Supply Adjustment of Older Individuals in the Long Run.” *Journal of Labor Economics*, 3(2), 237-253.

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.

Hausman, Jerry. 1981. “Labor Supply.” In Aaron, Henry and Pechman, Joseph., *eds*., How Taxes Affect Economic Behavior (pp. 27-71). Washington, DC: Brookings Institution.

Heckman, James. 1979. “Sample Selection Bias as a Specification Error.” *Econometrica*, 47(1), 153-161.

Kleven, Henrik, and Mazhar Waseem. 2013. “Using Notches to Uncover Optimization Frcitions and Structural Elasticities: Theory and Evidence from Pakistan.” *Quarterly Journal of Economics*, 128(2), 669-723.

Laitner, John, and Dan Silverman. 2012. “Consumption, Retirement, and Social Security: Evaluating the Efficiency of Reform that Encourages Longer Careers.” *Journal of Public Economics*, 96, 615-634.

Leonesio, Michael. 1990. “The Effects of the Social Security Earnings Test on the Labor-Market Activity of Older Americans: A Review of Empirical Evidence.” *Social Security Bulletin*, 53(5), 2-21.

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.

Meyer, Bruce, and Dan Rosenbaum. 2001. “Welfare, the Earned Income Tax Credit, and the Labor Supply of Single Mothers.” *Quarterly Journal of Economics*, 116(3), 1063-1114.

Powell, James. 1994. “Estimation of Semiparametric Models.” In Engle, R. F., and McFadden, D., *eds*., *Handbook of Econometrics* (pp. 2443-2521). Amsterdam: Elsevier.

Rubin, Donald. 1974. “Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies.” *Journal of Educational Psychology*, 66(5), 688-701.

Saez, Emmanuel. 2010. "Do Taxpayers Bunch at Kink Points?" *American Economic Journal: Economic Policy*, 2, 180-212.
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.
Tergesen, Anne. 2016. "How to Get People to Delay Retirement." *Wall St. Journal*, March 21.
Vroman, Wayne. 1985. "Some Economic Effects of the Retirement Test." In Ehrenberg, R., *ed*., Research in Labor Economics (Vol. 7, pp. 31-89). Greenwich, CT: Jai Press, Inc.

Figure 1. *Earnings Test Real Exempt Amount, 1981 to 1987*

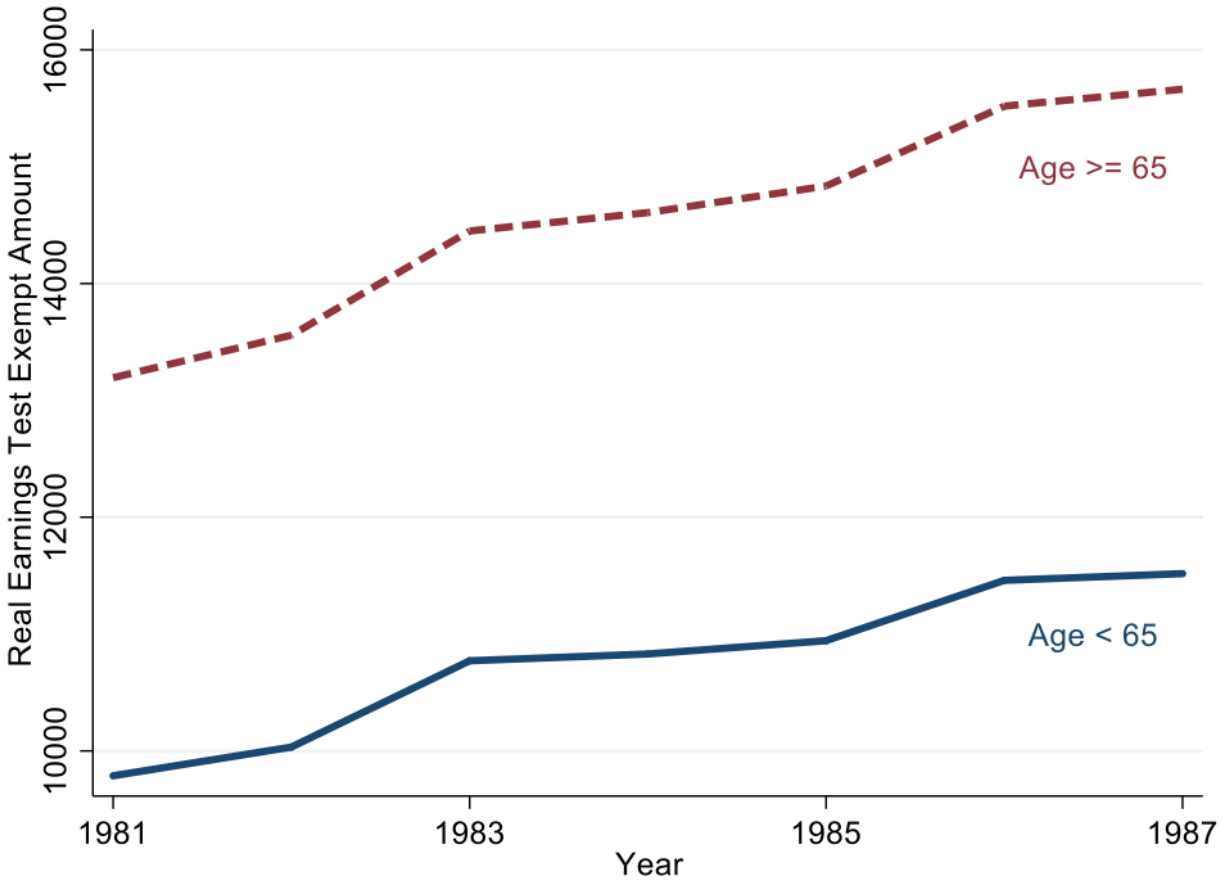

Notes: The figure shows the real value of the exempt amount over time among those 62-64 years old (labeled "Age<65" in the graph) and those 65 and above. The AET applied to earnings of claimants from ages 62 to 71 from 1981 to 1982, but only to claimants aged 62 to 69 from 1983 to 1989. All dollar figures are expressed in real 2010 dollars.

Figure 2. *Extensive Margin Incentives*

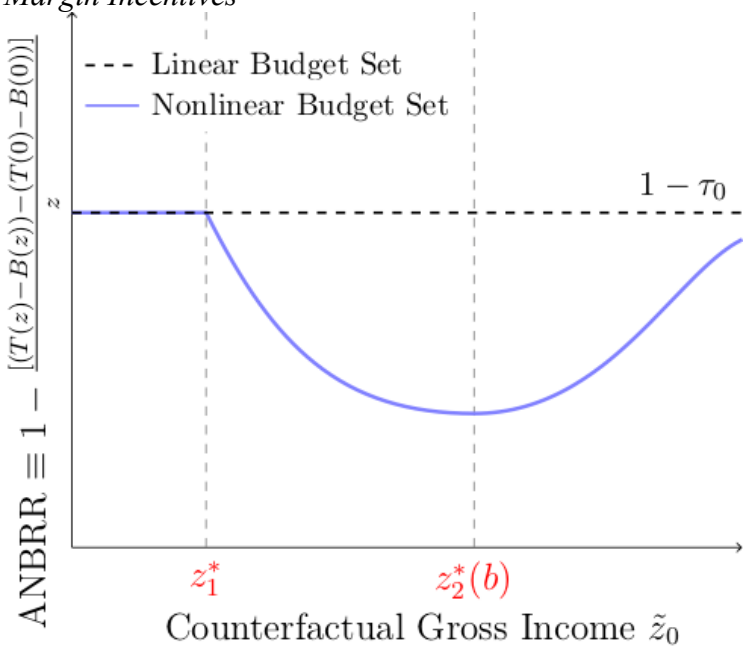

Notes: The figure shows the ANBRR (*y*-axis) as a function of gross income (*x*-axis). The ANBRR is defined as $ANBRR \equiv 1 - \frac{[(T(z)-B(z))-(T(0)-B(0))]}{z}$, *i.e.* the fraction of gross income an individual keeps, net of benefit reduction and taxes. Incentives under a linear tax schedule in which the ANBRR is equal to $1-\tau_0$ everywhere are represented by the dashed line. Incentives under a kinked tax schedule – in which the ANBRR is equal to $1-\tau_0$ below the first kink point z_1*, $1-\tau_0-\tau_b(z-z_1^*)/z$ between $z_1^*$and $z_2^*(b)$, and $1-\tau_0-b/z$ above $z_2^*(b)$ – are represented by the solid line.

Figure 3. *Extensive Margin Response by State 0 Counterfactual Earnings for a Kink*

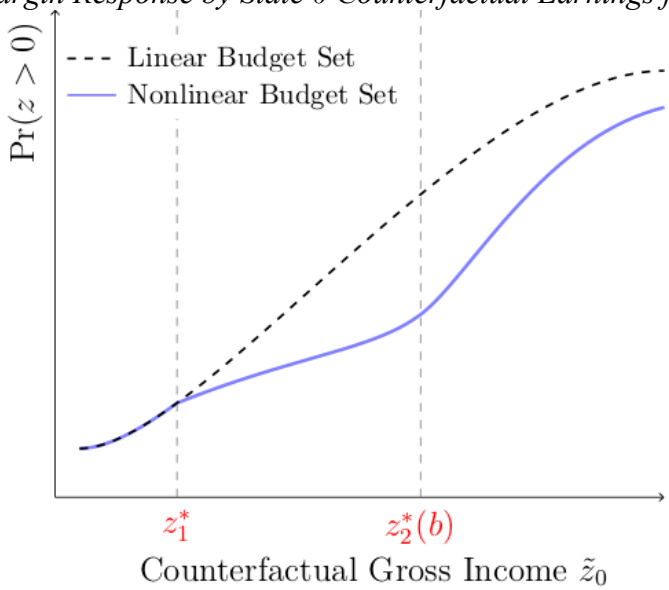

Notes: The figure illustrates the extensive margin response to the imposition of a kink, by counterfactual gross income. The *x*-axis shows desired gross income if employed on a linear budget set in state 0, *i.e.* $\tilde{z}_0$. The y-axis shows a hypothetical probability of employment, under two scenarios: a linear schedule in state 0 (dashed line) and a kinked tax schedule in state 1. The figure illustrates that the effect of the AET on employment grows as the AET reduces the ANBRR above the convex kink, and shrinks above the non-convex kink, eventually asymptoting to zero. If the marginal effect of the ANBRR on employment does not depend on counterfactual earnings, the magnitude of the effect of the AET on the employment rate should decrease beginning at the non-convex kink, which is the case shown in Figure 3. However, if the marginal effect of the ANBRR on employment depends on counterfactual earnings, then the effect of the AET on the employment rate may be maximized at a counterfactual earnings level other than the level associated with the non-convex kink. In either case, the effect of the AET on employment follows a U-shaped pattern, as in Figure 3.

Figure 4. *Histogram of Percent Real Earnings Growth, Ages 59 to 60*

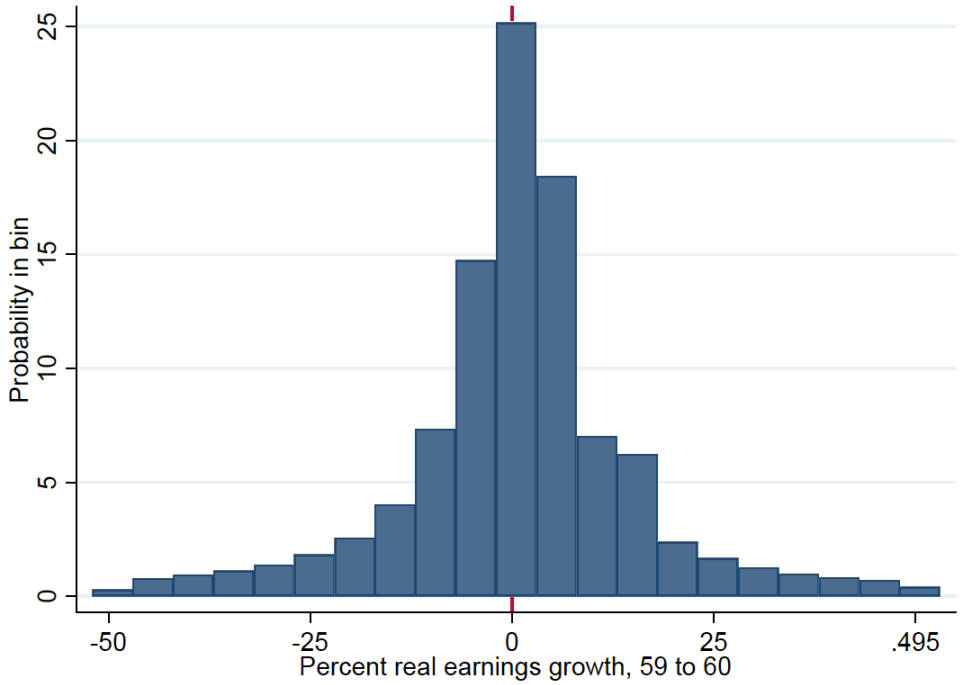

Notes: This histogram shows that there is a large mass at or near zero percent real earnings growth across a "placebo" set of ages, 59 and 60, when individuals do not face the AET. This indicates that a substantial mass of individuals have no growth in desired real earnings, consistent with the assumptions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

Figure 5. *Earnings Distributions at Ages 60 and 62*

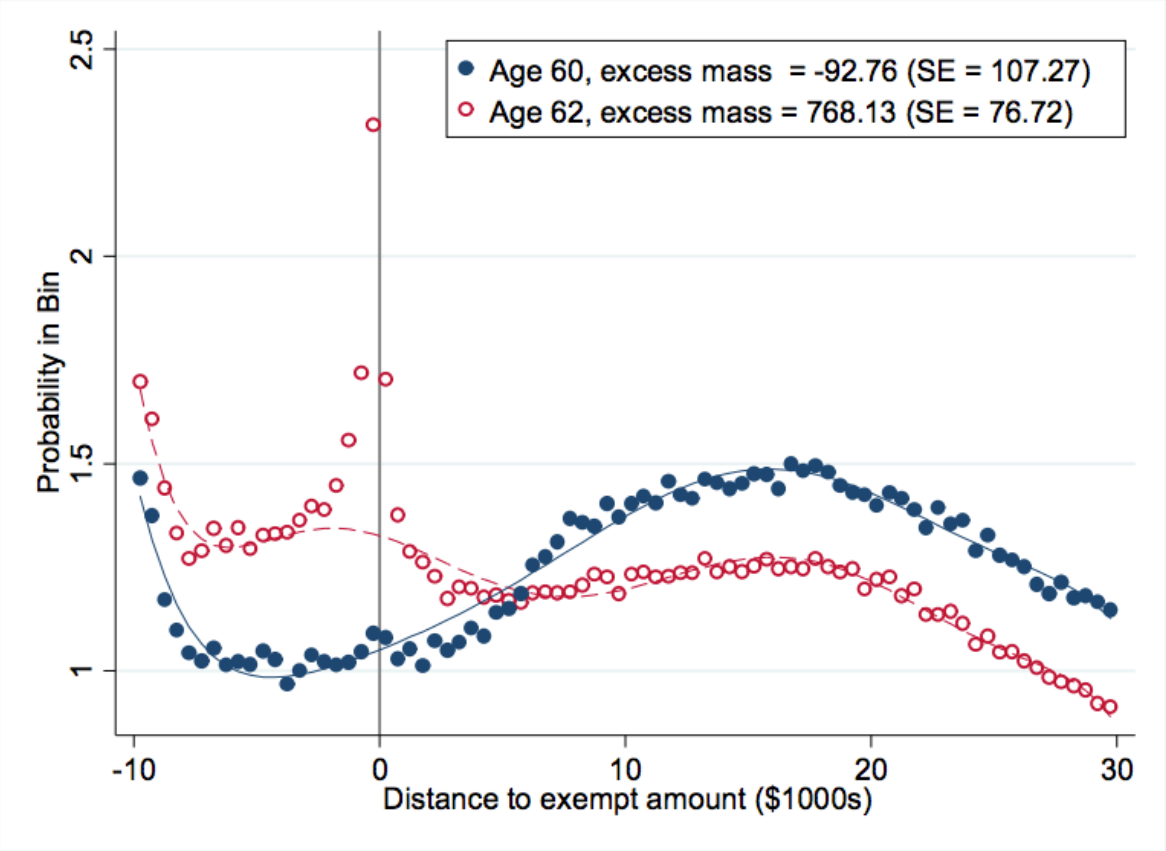

Note: the figure shows the earnings distributions at ages 60 and 62, based on a 25 percent sample of the SSA data. The sample consists of individuals with age 60 earnings that are positive and within $40,000 of the exempt amount, born 1918 to 1923, with no self-employment income at age 60, and excluding individuals who ever have negative earnings at ages 50-57 or 63-70. Source: Gelber *et al.* (2018).

Figure 6. *Probability of positive earnings by age and earnings relative to exempt amount*

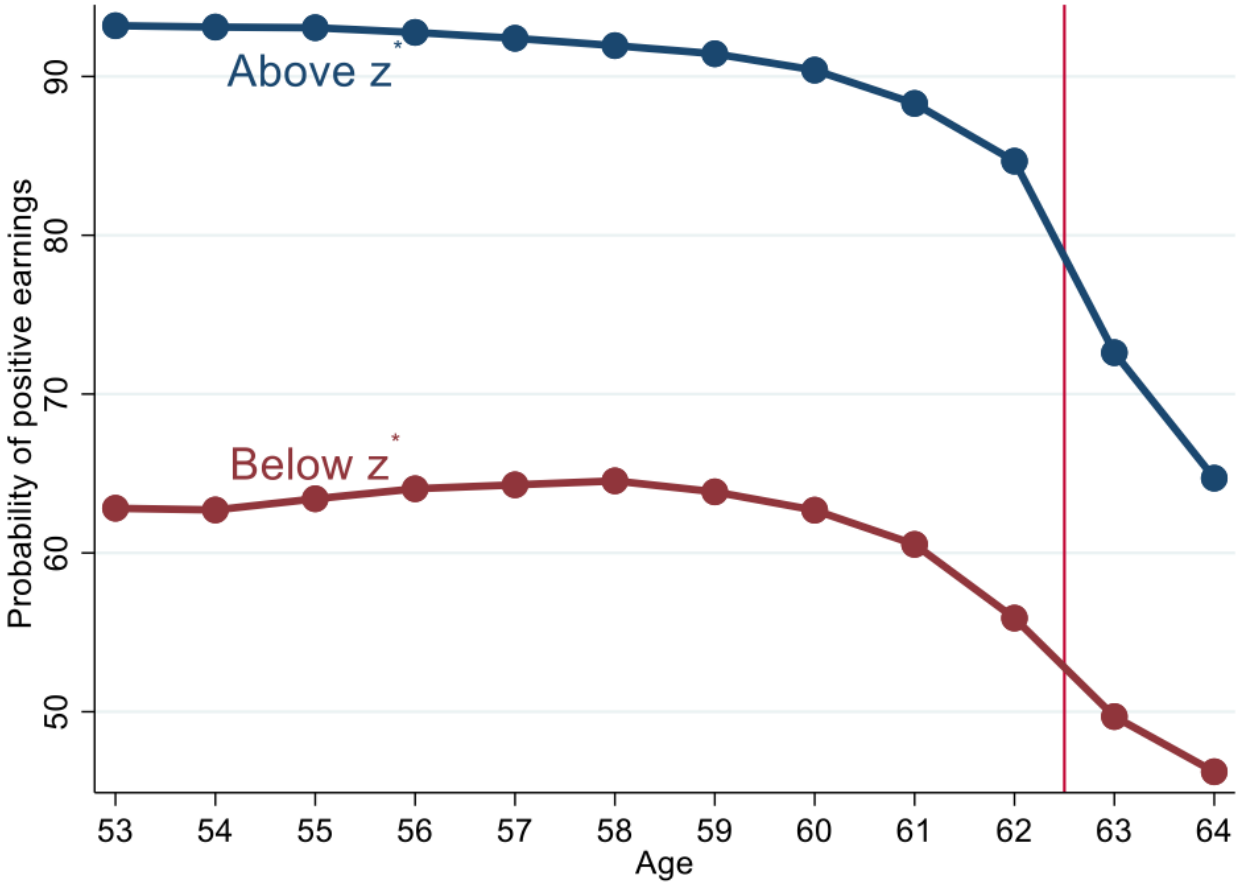

Notes: earnings relative to the exempt amount – *i.e.* below *vs.* above z^*—are measured at age a, whereas "age" on the x-axis and the employment probability on the y-axis are measured at age a+3. Source: SSA Administrative data and authors' calculations.

Figure 7. *Difference between probability of positive earnings among those earning above and below the exempt amount, by age*

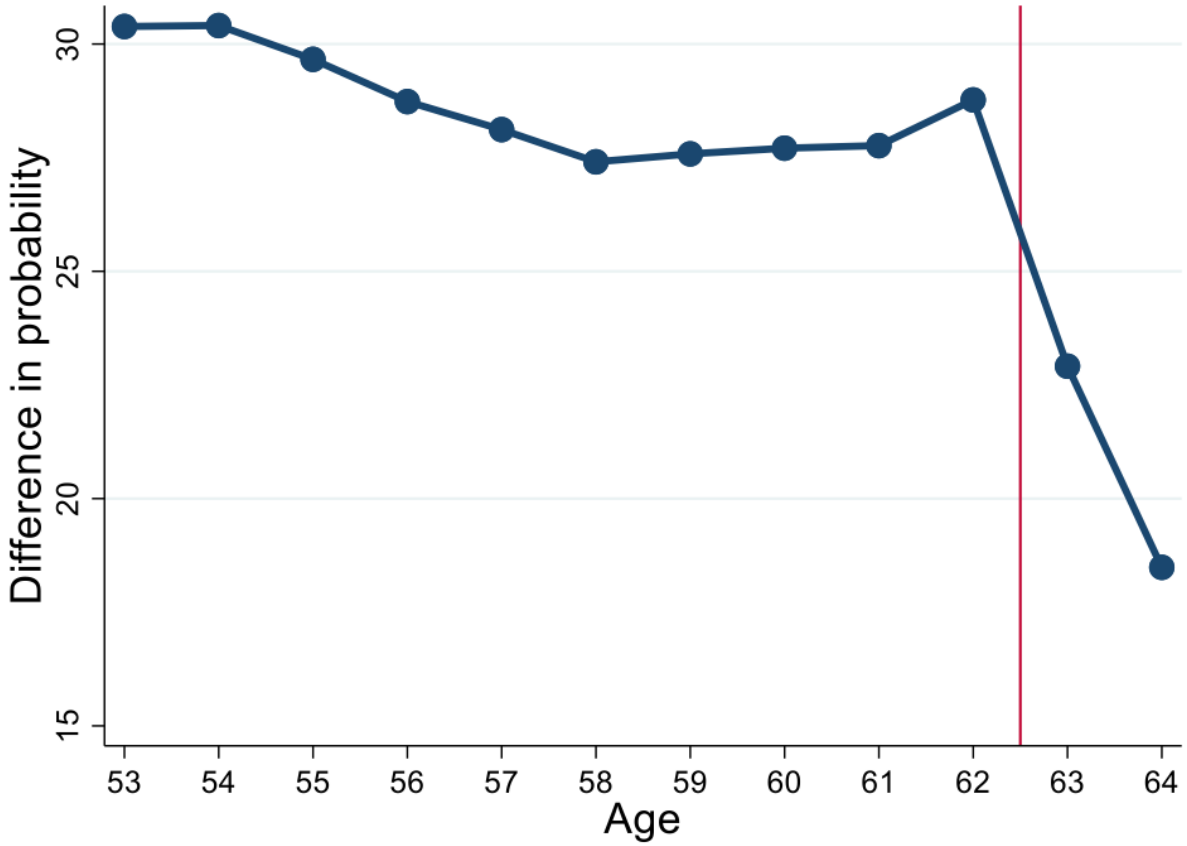

Notes: the figure shows (y-axis) the difference between those earning above z^* and those earning below z^* at age a in the probability of positive earnings at age a+3, as a function of age a+3 (x-axis). Source: SSA Administrative data and authors' calculations.

Figure 8. *Placebo estimates*

A. Placebo coefficient estimates (y-axis) using lower, placebo "exempt amounts" (x-axis, in real $2010)

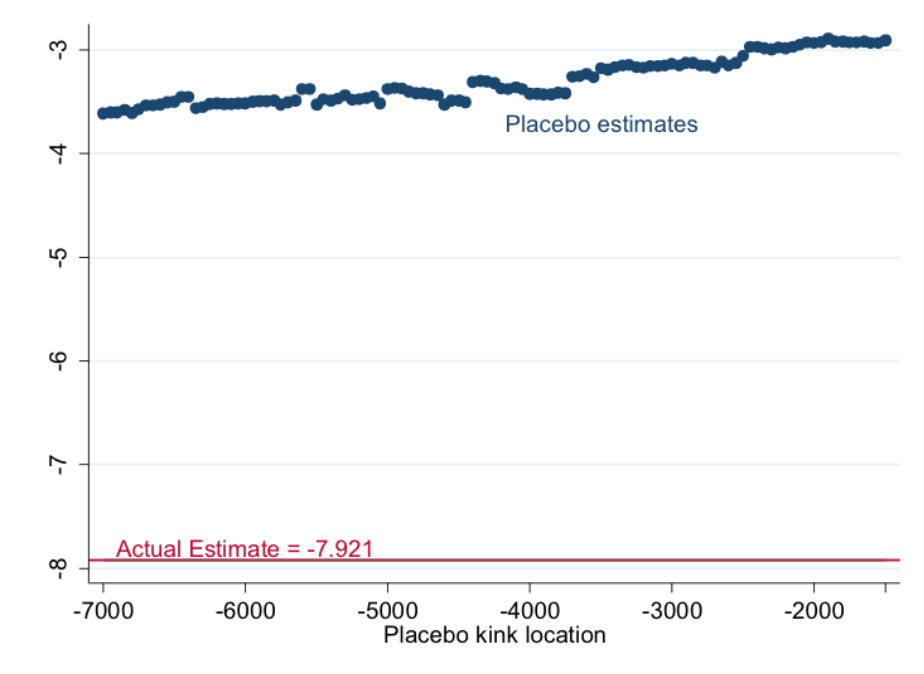

B. Placebo coefficient estimates (y-axis) at younger ages, using placebo "exempt amounts" (x-axis, in real $2010)

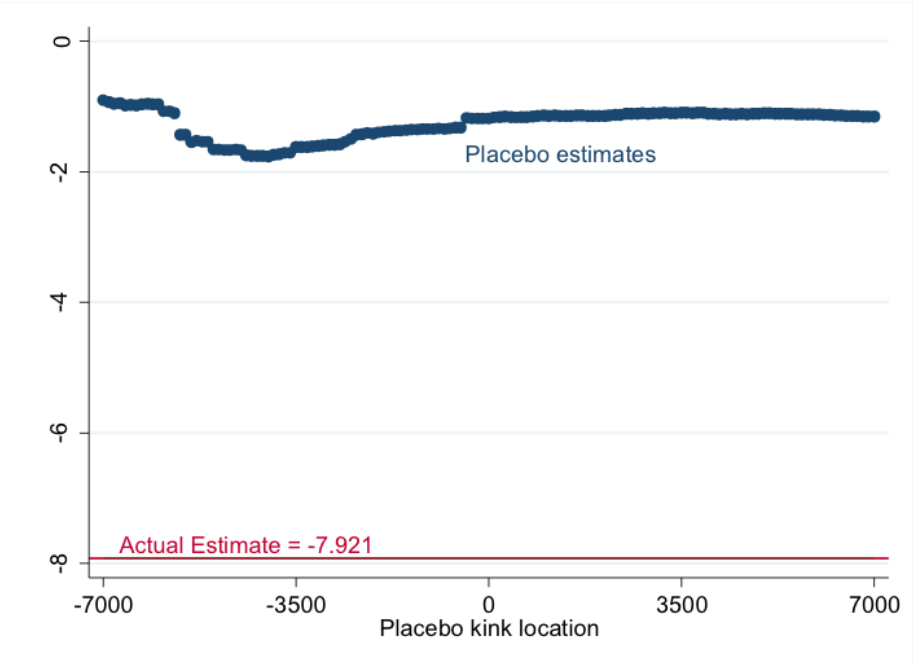

Source: SSA Administrative data and authors' calculations.

Figure 9. *Effect of AET on Employment, by Base Age Distance to the Exempt Amount*
by Base Age Distance to the Exempt Amount

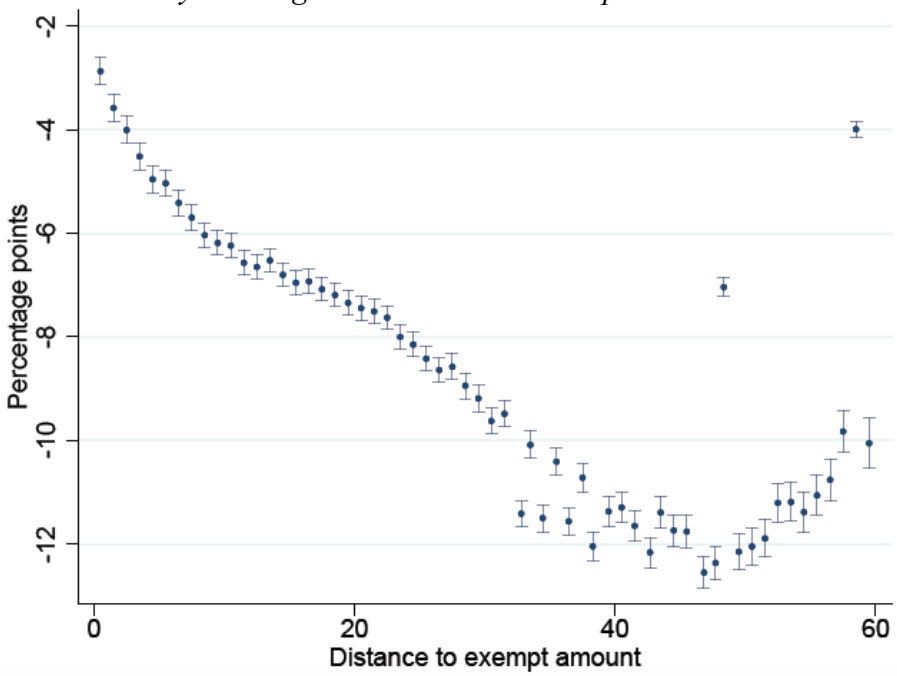

Notes: the figure shows the treatment effect of the AET on employment, for each bin in base year earnings relative to the exempt amount.

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

	Full Sample	Analysis Sample
Positive earnings dummy, ages 50-64	58.22	73.03
	(39.69)	(30.06)
Earnings, ages 50-64	20,396.08	25,617.26
	(21,987.37)	(21,769.64)
Claim Age	63.53	63.53
	(2.26)	(2.17)
Female dummy	45.55	42.99
	(49.80)	(49.51)
White dummy	88.11	89.50
	(32.37)	(30.66)
Year of birth	1920.55	1920.53
	(1.70)	(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=3*	(4) *t=4*
	Main		With controls	
Post x Treat	-7.066	-9.437	-7.157	-9.759
	(0.048)***	(0.047)***	(0.048)***	(0.047)***
# Obs	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

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 a ranging from a=55 to a=61. The controls in Columns 3 and 4 are a quadratic for age t. Here and in the other tables, *** denotes statistical significance of $p<0.01$ from a two-sided test of equality with zero. Standard errors here are clustered by individual; the estimates are also highly statistically significant when using permutation test described in the main text.

Table 3. *Estimates using ages 50-61*

	(1) *t=3*	(2) *t=4*	(3) *t=3*	(4) *t=4*
	Main		With controls	
Post x Treat	-7.921	-10.536	-7.872	-10.692
	(0.046)***	(0.043)***	(0.046)***	(0.043)***
# Obs	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

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 a=55 to a=61, whereas Table 3 uses the age range a=50 to a=61. This is the reason that the number of observations is larger in Table 3 than in Table 2; the number of people is also larger because we condition on positive earnings in the base year, and this is more prevalent on average in the sample in Table 3, which is younger on average and has higher average employment rates in the base year.

Table 4. *Controlling for age fixed effects*

	(1) *t+3*	(2) *t+4*	(3) *t+3*	(4) *t+4*
	t=50 to *t*=61		*t*=55 to *t*=61	
Post x	-7.176	-9.751	-7.988	-10.785
Treat	(0.048)***	(0.047)***	(0.046)***	(0.044)***

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 sample is the same as in Table 2.

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

	(1) Linear	(2) Cubic	(3) Decile dummies
Post x	-7.448	-5.705	-6.403
Treat	(0.195)***	(0.199)***	(0.199)***

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 the measure of OASI benefits we have in the SSA data – benefits received in the final year that the individual is recorded receiving OASI in the data – 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. See Appendix Table 1 for other notes.

Online Appendix

We briefly recapitulate the key findings of in Gelber *et al.* (2018), which further help to bolster the credibility of the results of the current paper. Paralleling the sharp change at the exempt amount in the slope of the ANBRR shown in Figure 2, Gelber *et al.* show theoretically that there should be a corresponding sharp change in slope of the employment rate as a function of age 60 earnings if there are frictions at the intensive margin that prevent individuals from adjusting to the AET by bunching at the exempt amount. This pattern does arise in the data: We show in Appendix Figure 1 (from Gelber *et al.*, 2018) that there is no visible change in the slope of the employment rate at the exempt amount at ages 61 and 62—prior to the ages when we should start to see an effect—but that we begin to observe a visible change in slope at ages 63 and 64. Using a Regression Kink Design (RKD), Gelber *et al.* (2018) show that there is no statistically significant change in slope at the exempt amount at ages 61 or 62, but that the change in slope becomes statistically significant at ages 63 and 64.

We also show in Appendix Figure 2 that predetermined covariates do not noticeably change in slope or level around the exempt amount, as the regressions in Gelber *et al.* (2018) confirm. This is consistent with the assumptions necessary for the validity of the empirical design.

Appendix Figure 1. *Probability of Positive Earnings by Single Year of Age, Ages 61 to 64*

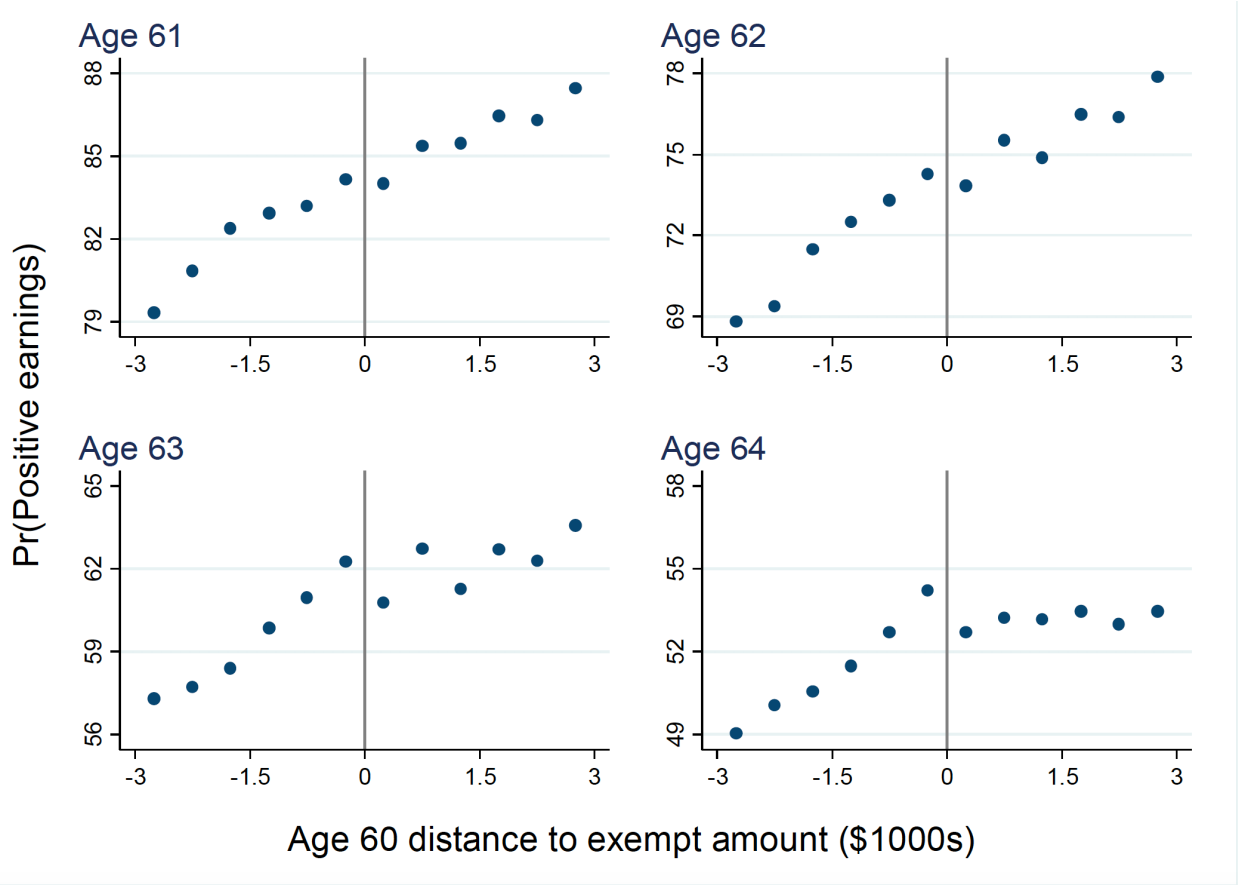

Notes: the source of the figure is Gelber *et al.* (2018). Each figure plots the mean annual employment rate, *i.e.* the probability of positive earnings, for each single year of age from 61 to 64, as a function of the distance to the exempt amount, which has been normalized to zero. The sample is individuals with positive age 60 earnings and no age 60 self-employment income, born 1918 to 1923.

Appendix Figure 2. *Predetermined covariates around the exempt amount*

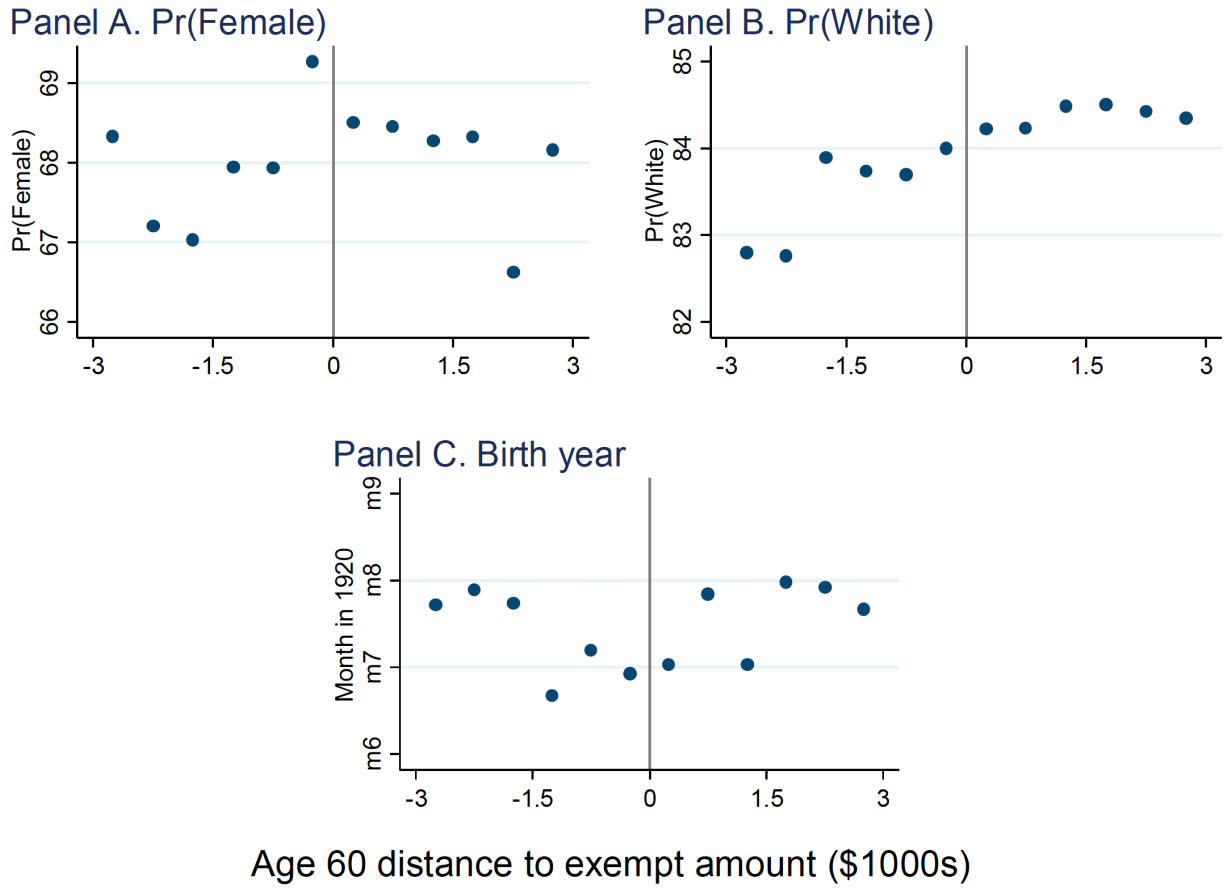

Notes: the source of the figure is Gelber *et al.* (2018). The figure shows the bin means of predetermined covariates as a function of the distance to the age 60 exempt amount. The figure demonstrates that there are no clear visual changes in slope in any of these covariates at the age 60 exempt amount, consistent with the assumptions necessary for the validity of the regression kink design employed in Gelber *et al.* (2018).

Appendix Table 1. *Results with additional controls*

	(1) *t=3*	(2) *t=4*	(3) *t=3*	(4) *t=4*	(5) *t=3*	(6) *t=4*
	W/ no controls		W/ demographic controls		W/ additional linear trends	
Post x	-8.127	-10.593	-7.182	-5.425	-8.176	-10.644
Treat	(0.186)***	(0.199)***	(0.290)***	(0.200)***	(0.290)***	(0.310)***
# Obs	2,853,073	2,853,073	2,853,073	2,853,073	2,853,073	2,853,073
# People	738,143	738,143	738,143	738,143	738,143	738,143

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 a ranging from a=55 to a=61. The controls in Columns 1 and 2 are dummies for female, white, and year of birth; Columns 3 and 4 use all of these controls and additionally add controls for linear trends in age that are separate in the treatment and control groups. The sample here is smaller than in the main text because due to data availability, in this table we use data beginning in 1978 (as opposed to the earlier tables, which use data beginning in 1968). However, this makes only a modest difference to the main estimates: Columns 1 and 2 of Appendix Table 1 show that when we run the main specification from Table 2 Columns 1 and 2 on the sample beginning in 1978, the results are similar to those shown in Table 2 in the main text (with standard errors that are correspondingly larger, but still implying statistically significant effects).