Online Appendix

A. Framework for interpreting the 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.¹ 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 received before accounting for the effects of the AET or taxes.² Current benefits, $B_j(z)$, are determined both by *b* as well as 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.³ This tax, which reduces net earnings relative to gross earnings, is separate from the AET, which only acts to reduce OASI benefits.⁴ 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:

¹ 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.

 $^{^{2}}$ 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.

³ 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. ⁴ 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.

$$B_{1}(z) = \begin{cases} b \text{ for } z \leq z_{1}^{*} \\ b - \tau_{b}(z - z_{1}^{*}) \text{ for } z_{1}^{*} < z \leq z_{2}^{*}(b) \\ 0 \text{ for } z_{2}^{*}(b) < z \end{cases}$$

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:

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

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

$$ANBRR \equiv 1 - \frac{\left[\left(T(z) - B(z) \right) - \left(T(0) - B(0) \right) \right]}{7}.$$

The ANBRR reflects the fraction of an individual's gross income that she keeps, net of both taxes and benefits, if she 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 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 .

⁵ Gelber *et al.* (2018) give an extended discussion of extensive margin earnings decisions in the presence of a kinked budget set.

B. Validating our Measure of Social Security Benefits

A key control variable in our analysis is the monthly benefit amount. Because we are interested in employment effects at ages 63 and 64, our interest is in what the benefit amounts would be if a person claimed OASI benefits. Our goal, therefore, is to impute the benefits a person would earn if she claimed at age 62. As d escribed in the text, we impute benefits using a calculator provided by the Social Security Administration, applied to observed earnings histories, but with earnings at ages 55-61 set equal to their age 55 level.

This procedure creates two potential sources of measurement error in our imputation. First, of course, we do not use actual earnings at ages 55-61. Second, we do not observe exact earnings; instead there is a slight amount of random rounding in principle leads to errors in the benefit imputation. To assess the severity of this measurement error, we use data from the BEPUF, in which we observe actual monthly benefits as of 2004. For most people in the BEPUF, benefits observed are not the object of interest, because most people are not 62 year-old claimants in 2004. However, for people who were born in 1942 and claim OASI benefits as a retired worker, actual benefits in the BEPUF are exactly the object of interest.

We validate our measure by estimating the following regression

$$b_i = \beta_0 + \beta_1 b_i^* + e_i$$

where *b* is our imputed benefit and b^* is the observed benefit, and We use data from the BEPUF and we limit the sample to people born in 1942 and claiming at age 62. The slope coefficient in this regression is the signal-to-noise ratio, and if \hat{b} were measured perfectly, we would expect a constant of 0 and a slope of 1. The results are in Appendix Table 1. The first column shows results for everyone and the second for retired workers only (e.g., excluding disabled beneficiaries). In the first column the slope is 0.844, close to 1 but clearly below it, and the constant is \$232. However, the sample in the first column includes people who claimed benefits on a spouse's record or as a disabled beneficiary, and our imputation is likely inaccurate for them. When we limit the sample to retired workers only in column (2), the slope coefficient rises to 0.95 and the constant falls to \$92. Only a small amount of measurement error remains. We conclude that our imputation process measures age 62 benefits faithfully.

C. Discussion of appendix figures

Appendix Figure 1 shows that 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 age a.⁶ 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.

Appendix Figure 1 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.

Appendix Figure 2 presents our main Figure 7 in a different way. Figure 7 shows treatment effects that are specific to bins of base age earnings. The treatment effect is the differential change in the probability of positive earnings, relative to the pre-period change and relative to the omitted bin. An alternative way to show this is to calculate, for each bin of base age distance to exempt amount and for each base age a, the probability of positive earnings at a + 3, and then to difference out this probability relative to some pre-period age. In Appendix Figure 2 we difference out the age 55 probabilities. The figure shows three important patterns. First, holding fixed earnings, as age gets higher, the probability of positive earnings falls. Second, at younger ages, the relationship between earnings and future employment is near zero, but, third, at higher ages, it is u-shaped: at first flat, then decreasing, then increasing. This is the basic pattern implied by our model.

We also 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 3 (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 63 and 64.

We also show in Appendix Figure 4 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.

⁶ 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 2018).

D. Interpretation of the results in light of benefit enhancement

One important question is whether the response to the AET is influenced by its impact on future OASI benefits, given that benefits may be enhanced in the future after they are initially reduced due to the AET. Several considerations point against the view that future benefits play a significant role in mediating responses to the Earnings Test, as we discuss in detail in this Appendix.

First, responses to the AET do not appear to be larger for those who have relatively short lifespans, for whom the Earnings Test is particularly punitive. As described above, literature has established that individuals "bunch" at the Earnings Test exempt amount (Friedberg 1998; Friedberg 2000; Gelber, Jones, and Sacks forthcoming). If expected future benefits matter for responses to the Earnings Test, then we would expect these bunchers to be disproportionately composed of those who have short average lifespans, for whom the Earnings Test is particularly punitive. If so, expected lifespan should fall sharply as a function of earnings, in a radius within approximately \$3,000 of the exempt amount where individuals tend to bunch (as shown in Gelber, Jones, and Sacks forthcoming). However, we do not see this pattern in the data. Appendix Figures 5 and 6 show that the probability of living past 70 (Appendix Figure 5) and average realized lifespan conditional on dying by the end of the sample (Appendix Figure 6), in a one percent sample of the SSA data on the U.S. population that was used in Gelber, Jones, and Sacks (forthcoming). Both graphs are essentially flat around the exempt amount, suggesting that future benefits do not play a role in mediating these responses to the AET.⁷

Second, the data show no response to the incentives created by future benefit enhancement. Over time benefit enhancement has become more generous but, as we describe in greater detail in Gelber, Jones, and Sacks (2019), for those over NRA there is no evidence of systematic bunching reaction to changes in the DRC and little relationship between bunching and life expectancy. For those under NRA, future benefits are enhanced substantially if current earnings exceed the exempt amount by even a single dollar. This creates an *upward* notch in lifetime income. Gelber, Jones and Sacks (2019) show however that there is disproportionate bunching *under* the exempt amount – the opposite of the pattern we would expect if people reacted to the full incentives created by benefit enhancement..

Third, in Table 3 of the current paper, we show heterogeneity in the estimated effects with respect to average lifetime income. Although income is not a perfect proxy for longevity, the two variables are significantly correlated (and as shown in literature from Preston 1975 to Chetty et al. 2016). Those with below-average income will on average have lower lifespan and, all else equal, should therefore react more to the Earnings Test because it is more punitive. In fact, we observe a smaller reaction among those with low average prior lifetime income.

Fourth, Table 3 also shows heterogeneity with respect to sex. Men have shorter average lifetimes than women, so we might expect to react more to the Earnings Test. In fact, women respond more than men.

Fifth, in Gelber, Jones, and Sacks (2019) we provide additional evidence that individuals mis-perceive the Earnings Test as more punitive than it is. Perhaps the reason for this apparent lack of reaction to variation in future benefits, is that the earnings test is widely mis-perceived as

⁷ We lack data on lifespan in the EPUF, and therefore we are not able to analyze the relationship between lifespan and extensive margin responses, given our current data constraints. However, the lack of a relationship between lifespan and intensive margin responses is telling and strongly suggestive that responses to the Earnings Test are not mediated by future benefits.

a pure tax. Most popular guides do not note the subsequent adjustment in benefits under the earnings test (Gruber & Orszag (2003)). During the period that we study, the popular guide Your Income Tax (J.K. Lasser Institute (1997)), for example, warned readers that if "you are under age 70, Social Security benefits are reduced by earned income," but did not note the subsequent benefit adjustment. 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.

E. Replication in the restricted-access data and in the BEPUF

We began this project with access to data from the Master Earning's File. The sample was a 25% random sample of people born 1918-1923. We limited that sample to people without self-employment income. The main advantage of the restricted-access data is that it contains exact earnings information, not top coded or rounded, and that it separates self-employment and other income; otherwise it is similar to the public access data sets. An additional advantage of the restricted data is its enormous size, although estimates with the public use files appear reasonably precise anyway.

Before we lost access to the restricted-access data, we estimated simple difference-indifferences models without controls. We present those results in Appendix Table 2. The coefficient when t = 3 is -7.1 and when t = 4 it is -9.4. These estimates are not comparable to our main results because of differences in the cohorts (1918-1923 vs. 1931-1943) and differences in the specifications. To make the estimates comparable, we present in Appendix Table 3 estimates from the EPUF that look at the same set of cohorts and, in columns (1) and (4), use the same set of controls. The estimates are quite similar, -8.2 and -8.6 when t = 3 and t = 4.

These results show very similar estimates in the restricted-access and public use files. The results do raise a different question, namely, why does responsiveness seem so much higher for the older cohorts? The answer is that the estimates in the restricted access data do not control for the confounding effect of benefits. When we control for imputed benefits in columns (2) and (5), the estimates fall dramatically and are now much closer to the estimates for later cohorts (although still 25 percent larger). Further controlling for demographics in the public use files does not much change the estimates. Thus we believe that we would obtain similar results in the restricted-access data, were we able to estimate our main specifications in them.

We also replicate our main results in the BEPUF. We follow our sample selection and specification as closely as possible. However, the BEPUF only contains earnings information through 2003, so we limit the sample to people born between 1931 and 1940 (rather than 1931 to 1943 in the EPUF, which runs through 2006), so that the last cohort in the BEPUF sample reaches age 63 at the end of the data. Because the BEPUF is a random sample of Social Security claimants, its sampling frame is quite different from the EPUF's. To make it more comparable, we further limit the BEPUF sample to people with retired worker benefits (as of 2004, the date at which benefits are recorded); otherwise the BEPUF would oversample disabled beneficiaries, who likely have different patterns of labor force attachment.

We present estimates from the BEPUF in Appendix Table 4. This table is exactly analogous to our Table 2 of the main text, our main results. The results are also highly similar. In our preferred specification, we estimate a DID coefficient of -2.5 when t = 3 and -2.8 when t = 4. Although the BEPUF's sampling frame is different from the EPUF's, we find it reassuring that both sets of estimates are within sampling error of each other. We conclude that our results are not sensitive to the choice data set.

Appendix References not cited in the main text

- Brown, Jeffrey, Arie Kapteyn, Olivia Mitchell, and Teryn Mattox. 2013. "Framing the Social Security Earnings Test." Wharton Pension Research Council Working Paper.
- Chetty, Raj, Michael Stepner, Sarah Abraham, Shelby Lin, Benjamin Scuderi, Nicholas Turner, Augustin Bergeron, and David Cutler. 2016. "The Association Between Income and Life Expectancy in the United States, 2001-2014." *JAMA* 315(16): 1750-1766.
- 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.
- Fetter, Daniel, and Lee Lockwood. 2018 "Government Old-Age Support and Labor Supply: Evidence from the Old Age Assistance Program." *American Economic Review* 108(8): 2174-2211.
- Hausman, Jerry. 1981. "Labor Supply." In *How Taxes Affect Economic Behavior*, ed. Henry Aaron and Joseph Pechman, 27-71. Washington, DC: Brookings Institution.
- Kleven, Henrik, and Mazhar Waseem. 2013. "Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan." *Quarterly Journal of Economics* 128(2): 669-723.
- Maestas, Nichole. 2010. "Back to Work: Expectations and Realizations of Work After Retirement." *Journal of Human Resources* 45(3): 718-748.
- Preston, Samuel H. 1975. "The Changing Relation between Mortality and Level of Economic Development." *Population Studies* 29(2): 231-248.
- 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(3): 180-212.



Appendix Figure 1. *Probability of positive earnings by age and earnings relative to exempt amount*

Notes: We show the employment probability by age, for the below-z* (above-z*) group that earned below (above) the exempt amount three years prior. To parallel our main specification, we plot means adjusting for our controls (female, age and cohort fixed effects, imputed benefits, and imputed benefits interacted with post). The shared area represents the 95% confidence region.

Appendix Figure 2 Differential probability of positive earnings at a + 3, relative to age 55, as a function of base age and base age earnings



Notes: Figure plots, for each base age a and bin of earnings relative to the exempt amount, the difference in the fraction of observations with positive earnings at a + 3, relative to the fraction in base age 55. The sample is drawn from the EPUF and consists of people born 1931-1943, with positive earnings in the indicated base ages.



Appendix Figure 3. 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 4. 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 Figure 5. Probability of living to age 70 as a function of age 62-64 earnings

Notes: The sample consists of people aged 62-64 in 1990-1999 who claimed by age 65. The x-axis is earnings relative to the exempt amount. The y-axis shows the fraction of people living to age 70 or greater in each \$800 bin. The shaded area is the 95% confidence interval. Source: SSA data.



Appendix Figure 6. Age at death as a function of age 62-64 earnings

Notes: The sample consists of people aged 62-64 in 1990-1999 who claimed by age 65. The x-axis is earnings relative to the exempt amount. The y-axis shows the average age at death (conditional on dying) in each \$800 bin. The shaded area is the 95% confidence interval. Source: SSA data.

Sample	All	Retired workers only
	(1)	(2)
Slope	0.844 (0.005)	0.953 (0.004)
Constant	231.888 (6.215)	91.679 (4.848)
R ² # Observations	$0.822 \\ 8,504$	0.934 7,372

Appendix Table 1. Validating benefit imputation

Notes: Table reports the estimates from a bivariate regression of monthly benefit amounts as reported in the BEPUF against imputed benefit amounts. Sample is limited to people born in 1942 and claiming benefits at age 62 because for this sample we observe actual benefits at age 62 in the BEPUF (i.e. the variable we impute). Robust standard errors in parentheses.

	t = 3 years ahead (1)	t = 4 years ahead (2)
DID coefficient	-7.1	-9.4
	(0.05)	(0.05)
# Observations	48,580,452	48,580,452
# People	8,296,628	8,296,628

Appendix Table 2. *Results from restricted-access administrative data and specifications that do not control for benefits*

Notes: Table shows the DID coefficient. The outcome is a dummy for positive earnings (x100) in the indicated number of years ahead. The sample is a 25% random sample of people born 1918-1923 from the Master Earnings File. The sample is limited to observations with positive earnings and base age 55 to 61. Additional controls include *post* and *treat*. Standard errors (in parentheses) are clustered by individual.

	t = 3 years ahead			t = 4 years ahead		
	(1)	(2)	(3)	(4)	(5)	(6)
DID Coefficient	-8.2 (0.4)	-3.8 (0.47)	-3.8 (0.47)	-8.6 (0.4)	-4.1 (0.5)	-4.1 (0.5)
#Observations #People	542,581 97,969	542,581 97,969	542,581 97,969	475,796 96,999	475,796 96,999	475,796 96,999
Controls						
Benefits		Х	Х		Х	Х
Sex			Х			Х
Age FE			Х			Х
Cohort FE			Х			Х

Appendix Table 3. Limiting the EPUF sample to earlier cohorts to match SSA data

Notes: Table shows the DID coefficient. The outcome is a dummy for positive earnings (x100) in the indicated number of years ahead. The sample is drawn from the EPUF. It consists of people born 1918-1923 with positive base age earnings, and base age 55 to 64 - t. All specifications control for main effects of *post* and *treat*. The benefit controls are imputed benefits plus their interaction with *post*. Standard errors (in parentheses) are clustered by individual.

	t = 3 years ahead			t = 4 years ahead		
	(1)	(2)	(3)	(4)	(5)	(6)
DID Coefficient	-3.5 (0.3)	-2.4 (0.3)	-2.5 (0.4)	-3.9 (0.4)	-2.7 (0.4)	-2.8 (0.4)
Elasticity	0.24	0.19	0.19	0.27	0.22	0.22
#Observations #People	772.417 133,240	772.417 133,240	772.417 133,240	672,596 132,093	672,596 132,093	672,596 132,093
Controls Benefits Sex Age FE Cohort FE		Х	X X X X		Х	X X X X

Appendix Table 4. *Replicating the main estimates in the BEPUF*

Notes: Table shows the DID coefficient. The outcome is a dummy for positive earnings (x100) in the indicated number of years ahead. The sample is drawn from the BEPUF. It consists of people born 1931-1940 with positive base age earnings, and base age 55 to 64 - t. All specifications control for main effects of *post* and *treat*. The benefit controls are imputed benefits plus their interaction with *post*. The elasticity assumes a 40% OASI claiming rate. Standard errors (in parentheses) are clustered by individual.